



HAL
open science

A PAC-Bayesian Link Between Generalisation and Flat Minima

Maxime Haddouche, Paul Viillard, Umut Şimşekli, Benjamin Guedj

► **To cite this version:**

Maxime Haddouche, Paul Viillard, Umut Şimşekli, Benjamin Guedj. A PAC-Bayesian Link Between Generalisation and Flat Minima. ALT 2025 - 36th International Conference on Algorithmic Learning Theory, Feb 2025, Milan, Italy. pp.1-31. hal-04455639v2

HAL Id: hal-04455639

<https://hal.science/hal-04455639v2>

Submitted on 12 Feb 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial - ShareAlike 4.0 International License

A PAC-Bayesian Link Between Generalisation and Flat Minima

Maxime Haddouche

*Inria - Département d'Informatique de l'Ecole Normale Supérieure
PSL Research University
Paris, France*

MAXIME.HADDOUCHE@INRIA.FR

Paul Viallard

*Univ. Rennes, Inria, CNRS IRISA - UMR 6074
Rennes, France*

PAUL.VIALLARD@INRIA.FR

Umut Şimşekli

*Inria - Département d'Informatique de l'Ecole Normale Supérieure
PSL Research University
Paris, France*

UMUT.SIMSEKLI@INRIA.FR

Benjamin Guedj

*Inria and University College London
London, UK*

BENJAMIN.GUEDJ@INRIA.FR

Editors: Gautam Kamath and Po-Ling Loh

Abstract

Modern machine learning usually involves predictors in the overparameterised setting (number of trained parameters greater than dataset size), and their training yields not only good performance on training data, but also good generalisation capacity. This phenomenon challenges many theoretical results, and remains an open problem. To reach a better understanding, we provide novel generalisation bounds involving gradient terms. To do so, we combine the PAC-Bayes toolbox with Poincaré and Log-Sobolev inequalities, avoiding an explicit dependency on the dimension of the predictor space. Our results highlight the positive influence of *flat minima* (being minima with a neighbourhood nearly minimising the learning problem as well) on generalisation performance, involving directly the benefits of the optimisation phase.

Keywords: Generalisation Bounds, PAC-Bayes, Flat Minima, Poincaré, Log-Sobolev Inequalities

1. Introduction

Understanding generalisation in modern machine learning problems has been a major challenge in learning theory. The goal here is to upper-bound the so-called *generalisation error*, which is the gap between the population and empirical risks, $R_{\mathcal{D}}(h) - \hat{R}_{\mathcal{S}_m}(h)$. Here, $h \in \mathbb{R}^d$ represents the parameters of a predictor, $R_{\mathcal{D}}(h) := \mathbb{E}_{\mathbf{z} \sim \mathcal{D}}[\ell(h, \mathbf{z})]$ denotes the population risk, \mathcal{D} is an unknown data distribution, ℓ is a loss function, $\hat{R}_{\mathcal{S}_m}(h) := \frac{1}{m} \sum_{i=1}^m \ell(h, \mathbf{z}_i)$ is the empirical risk, and $\mathcal{S}_m := \{\mathbf{z}_1, \dots, \mathbf{z}_m\}$ is a dataset in which each \mathbf{z}_i is independent and identically distributed (*i.i.d.*) with respect to \mathcal{D} .

Dating back to [Hochreiter and Schmidhuber \(1997\)](#), it has been hypothesised that the notion of ‘flatness’ (or sometimes equivalently referred to as ‘sharpness’) is closely linked to the generalisation error: among the minima found by the learning algorithm, the flatter the minimum, the lower the generalisation error. While the initial concept of flatness was (vaguely) defined through low Kolmogorov complexity, there is no globally accepted definition of flatness. Therefore, several notions

of flatness have been considered, typically based on the second-order derivatives of the empirical risk around the local minimum found by the learning algorithm, such as $\text{trace}(\nabla_h^2 \hat{R}_{\mathcal{S}_m}(h))$ (see *e.g.* Jastrzebski et al., 2017; Wen et al., 2023).

While there have been several attempts to link some form of flatness to generalisation in a mathematically rigorous manner (Neyshabur et al., 2017; Petzka et al., 2021; Yue et al., 2023; Andriushchenko et al., 2023), mainly with the framework of ‘sharpness aware minimisation’ (Foret et al., 2021), it has been shown recently that flat minima do not always imply good generalisation. In fact, there are scenarios such that the flattest minima result in the worst generalisation performance compared to non-flat ones (Wen et al., 2023).

In this study, we aim to develop novel links between flatness and generalisation error from a PAC-Bayesian perspective (Shawe-Taylor and Williamson, 1997; McAllester, 1999); see also (Guedj, 2019; Hellström et al., 2023; Alquier, 2024) for an introduction. While existing work involving PAC-Bayes and flatness either uses PAC-Bayes bounds as complexity measures (Neyshabur et al., 2017; Jiang et al., 2020; Dziugaite et al., 2020; Viallard et al., 2024a) or derives bounds for specific algorithms (Negrea et al., 2019; Neu, 2021), our aim is to derive general PAC-Bayes bounds involving flatness through gradient norm. Denoting by Q , the probability distribution of the algorithm’s output h , we identify sufficient conditions on Q such that flatness always implies good generalisation. More precisely, we make the following contributions:

- We show that, when Q satisfies the Poincaré inequality and a technical condition that we identify, we can obtain a ‘time-uniform estimation’ PAC-Bayes bound that mainly contains two terms:
 - (i) The flatness term: either $\mathbb{E}_{\mathbf{z} \sim \mathcal{D}} \mathbb{E}_{h \sim Q} [\|\nabla_h \ell(h, \mathbf{z})\|^2]$ or $\mathbb{E}_{h \sim Q} [\frac{1}{m} \sum_{i=1}^m \|\nabla_h \ell(h, \mathbf{z}_i)\|^2]$. The latter is directly linked to the Hessian of the loss ℓ , due to the connection between the Fisher information and the Hessian of the loss (Bickel and Doksum, 2015). For instance, under certain conditions, it can be shown that the trace is linked to gradients as $\text{trace}(\nabla_h^2 \hat{R}_{\mathcal{S}_m}(h)) = \frac{2}{m} \sum_{i=1}^m \|\nabla_h \ell(h, \mathbf{z}_i)\|^2$ (Wen et al., 2023, Lemma 4.1).
 - (ii) The classical PAC-Bayesian complexity term $\text{KL}(Q, P)$, where KL denotes the Kullback-Leibler divergence and P is data-independent ‘prior’ distribution.
- We further analyse the term $\text{KL}(Q, P)$. We show that when Q is a Gibbs distribution, *i.e.* $dQ(h) \propto \exp(-\gamma \hat{R}_{\mathcal{S}_m}(h)) dP(h)$ for some $\gamma > 0$ and P satisfies a log-Sobolev inequality, the generalisation error can be controlled *solely* by the term: $\gamma^2 c_{LS}(P) \mathbb{E}_{h \sim Q} [\|\nabla_h \hat{R}_{\mathcal{S}_m}(h)\|^2]$, where $c_{LS}(P)$ denotes the log-Sobolev constant of the prior P .
- We go beyond KL divergence to link flat minima to deterministic predictors (*i.e.* when Q is a Dirac distribution) through a novel Wasserstein-based generalisation bound for gradient Lipschitz loss functions.

We provide a numerical assessment of the technical condition underlying our main result, suggesting that it is suitable for neural networks on classification tasks, confirming the relevance of our bounds to better understand the generalisation ability of such models. Our results further shed light on the impact of the flatness of minima on generalisation error: when the learning algorithm ensures a sufficiently regular distribution over the parameters, the generalisation error can be directly controlled by the flatness of the region found by the algorithm.

2. Preliminaries

Framework. We consider a predictor set $\mathcal{H} \subseteq \mathbb{R}^d$ equipped with a norm $\|\cdot\|$, a data space \mathcal{Z} and the space of distributions $\mathcal{M}(\mathcal{H})$ over \mathcal{H} . We also consider a loss function $\ell : \mathcal{H} \times \mathcal{Z} \rightarrow \mathbb{R}$. We assume that we have an *i.i.d.* dataset $\mathcal{S} = (\mathbf{z}_i)_{i \geq 1} \in \mathcal{Z}^{\mathbb{N}}$ with associated unknown data distribution \mathcal{D} . For each $m \geq 1$, we define $\mathcal{S}_m := \{\mathbf{z}_1, \dots, \mathbf{z}_m\}$. In PAC-Bayes learning, we construct a data-driven posterior distribution $Q \in \mathcal{M}(\mathcal{H})$ with respect to a prior distribution P . To assess the generalisation ability of a predictor $h \in \mathcal{H}$, we define the *population risk* to be $R_{\mathcal{D}}(h) := \mathbb{E}_{\mathbf{z} \sim \mathcal{D}}[\ell(h, \mathbf{z})]$ and for each m , its empirical counterpart as $\hat{R}_{\mathcal{S}_m}(h) := \frac{1}{m} \sum_{i=1}^m \ell(h, \mathbf{z}_i)$. We also define the expected and empirical risks for $Q \in \mathcal{M}(\mathcal{H})$ as $R_{\mathcal{D}}(Q) := \mathbb{E}_{h \sim Q}[R_{\mathcal{D}}(h)]$ and $\hat{R}_{\mathcal{S}_m}(Q) := \mathbb{E}_{h \sim Q}[\hat{R}_{\mathcal{S}_m}(h)]$. PAC-Bayes bounds usually aim to control the *expected generalisation error* for each dataset size m , *i.e.* $R_{\mathcal{D}}(Q) - \hat{R}_{\mathcal{S}_m}(Q)$.

Background on Poincaré and log-Sobolev inequalities. In this work, we exploit Poincaré and log-Sobolev inequalities in the PAC-Bayes framework. We first recall their definitions: for a fixed distribution $Q \in \mathcal{M}(\mathcal{H})$, we define the *Sobolev space of order 1* on \mathbb{R}^d as follows:

$$H^1(Q) := \left\{ f \in L^2(Q) \cap D_1(\mathbb{R}^d) \mid \|\nabla f\| \in L^2(Q) \right\},$$

where $D_1(\mathbb{R}^d)$ denotes the set of derivable functions $f : \mathbb{R}^d \rightarrow \mathbb{R}$ and $L^2(Q)$ is the space of square-integrable function. In other words, $H^1(Q)$ is the set of functions that are square-integrable, with their gradient's norm also being square-integrable.

Definition 1 (Poincaré inequality) *A distribution Q satisfies a Poincaré inequality with constant $c_P(Q)$ if for all function $f \in H^1(Q)$ we have*

$$\text{Var}_{h \sim Q}(f(h)) \leq c_P(Q) \mathbb{E}_{h \sim Q} [\|\nabla_h f(h)\|^2],$$

where $\text{Var}_{h \sim Q}(f(h)) = \mathbb{E}_{h \sim Q}[f(h) - \mathbb{E}_{h \sim Q}[f(h)]]^2$ is the variance of f with respect to Q . We then say that Q is *Poincaré* with constant $c_P(Q)$, or that Q is *Poinc*(c_P).

Definition 2 (Log-Sobolev inequality) *A distribution Q satisfies a log-Sobolev inequality with constant $c_{LS}(Q)$ if for all function $f \in H^1(Q)$ we have*

$$\text{Ent}_{h \sim Q}(f^2(h)) := \mathbb{E}_{h \sim Q} \left[f^2(h) \log \left(\frac{f^2(h)}{\mathbb{E}_{h \sim Q}[f^2(h)]} \right) \right] \leq c_{LS}(Q) \mathbb{E}_{h \sim Q} [\|\nabla_h f(h)\|^2],$$

where the term $\text{Ent}_{h \sim Q}(f^2(h))$ is the entropy of f^2 . We then say that Q is *log-Sobolev* with constant $c_{LS}(Q)$, or that Q is *L-Sob*(c_{LS}).

The class of Gaussian distributions is an important particular case of distributions satisfying both Poincaré and log-Sobolev inequalities; this is the subject of Proposition 3.

Proposition 3 (Gross (1975); Brascamp and Lieb (1976); Beckner (1989)) *Given a distribution $Q = \mathcal{N}(\mu, \Sigma)$, where μ is the mean and Σ is the covariance matrix in \mathbb{R}^d . Then, for any $f \in H^1(Q)$:*

$$\text{Ent}_{h \sim Q}(f^2(h)) \leq 2 \mathbb{E}_{h \sim Q} [(\Sigma \nabla_h f(h), \nabla_h f(h))], \quad \text{and} \quad \text{Var}_{h \sim Q}(f(h)) \leq \mathbb{E}_{h \sim Q} [(\Sigma \nabla_h f(h), \nabla_h f(h))].$$

Thus, the distribution Q is *L-Sob*(c_{LS}) with constant $c_{LS}(Q) = 2\|\Sigma\|_{op}$ and is also *Poinc*(c_{LS}) with constant $c_{LS}(Q) = \|\Sigma\|_{op}$, where $\|\cdot\|_{op}$ denotes the operator norm.

In Proposition 3, the first inequality can be derived from the classical log-Sobolev inequality for $\mathcal{N}(\mathbf{0}, \text{Id})$ stated in Gross (1975), with a change of variable. Similarly, the Poincaré inequality can be obtained through a change of variable from the Poincaré inequality for $\mathcal{N}(\mathbf{0}, \text{Id})$ which is a particular case of the Brascamp-Lieb inequality for log-concave probability measures (Brascamp and Lieb, 1976) and is stated explicitly in Beckner (1989, Theorem 1).

We now focus on specific posterior distributions called *Gibbs posteriors*, or *Gibbs distributions*. Given a fixed loss $\ell : \mathcal{H} \times \mathcal{Z} \rightarrow \mathbb{R}$, a prior $P \in \mathcal{M}(\mathcal{H})$ and a dataset \mathcal{S}_m , the Gibbs posterior $Q_{\mathcal{S}_m}^\gamma$ is defined such that $dQ_{\mathcal{S}_m}^\gamma(h) \propto \exp(-\gamma \hat{R}_{\mathcal{S}_m}(h)) dP(h)$, where $\gamma > 0$ is an *inverse temperature*. Gibbs posteriors are a class of closed-form solutions for relaxation of Catoni (2007, Theorem 1.2.6) stated, for instance, in Alquier et al. (2016, Theorem 4.1). Proposition 4 shows that when the prior and the loss satisfy a few properties, then the associated Gibbs posterior is $L\text{-Sob}(c_{LS})$.

Proposition 4 *Assume that P is a probability measure on \mathcal{H} such that $dP(h) \propto \exp(-V(h))$ with $V : \mathcal{H} \rightarrow \mathbb{R}$ a smooth function such that $\text{Hess}(V) \succeq \frac{2}{c_{LS}(P)} \text{Id}$.¹ Assume that $\ell(h, \mathbf{z}) = \ell_1(h, \mathbf{z}) + \ell_2(h, \mathbf{z})$ with ℓ_1 convex, twice differentiable and ℓ_2 bounded. Then for any $\gamma > 0$, the Gibbs posterior $Q_{\mathcal{S}_m}^\gamma$ is $L\text{-Sob}(c_{LS})$ with constant $c_{LS}(Q_{\mathcal{S}_m}^\gamma) = c_{LS}(P) \exp(4\|\ell_2\|_\infty)$.*

Proposition 4 applies, e.g. when P is a Gaussian prior $P = \mathcal{N}(\mu_P, \Sigma_P)$. Notice that in this case $c_{LS}(P) = 2\|\Sigma_P\|_{op}$. This property is a straightforward application of Chafaï (2004, Corollary 2.1) with Guionnet and Zegarliński (2003, Property 2.6) and is stated in Appendix A for completeness. Finally, notice that satisfying a log-Sobolev inequality is stronger than satisfying a Poincaré one. This is stated for instance in Ledoux (2006, Proposition 2.1) and properly recalled in Appendix A.

3. Reaching a flat minimum allows Poincaré posteriors to generalise well

In this section, we consider posterior distributions Q that are $\text{Poinc}(c_P)$. This assumption covers the important case of Gaussian measures (Proposition 3) as well as all measures satisfying a log-Sobolev inequality (Proposition 15).

3.1. Time-uniform estimation PAC-Bayes bounds for heavy-tailed losses

We now focus on *time-uniform estimation* PAC-Bayes bounds, i.e. bounds such that there exists $\mathcal{C} \subseteq \mathcal{M}(\mathcal{H})$, and $\alpha > 0$ with probability at least $1 - \delta$, for all $Q \in \mathcal{C}$, and $m > 0$, there exists $\varepsilon_m > 0$ s.t.

$$R_{\mathcal{D}}(Q) \leq \frac{\alpha}{m} [\text{KL}(Q, P) + \log(1/\delta)] + \varepsilon_m.$$

Here, 'time-uniform' means that the bound holds with probability $1 - \delta$ for all m , and 'estimation' means that we directly control $R_{\mathcal{D}}$ instead of the generalisation gap $R_{\mathcal{D}} - \hat{R}_{\mathcal{S}_m}$. In particular, if $\sup_{m \geq m_0} \varepsilon_m \leq \varepsilon$ for some small $\varepsilon > 0$, we interpret a time-uniform estimation bound as a *transitory fast rate*, i.e. a bound decaying for all $m \geq m_0$ below 2ε at speed $1/m$. Such a property is of interest to understand why deep neural networks rapidly acquire a good generalisation ability. Another fundamental difference between time-uniform PAC-Bayes bounds (which recently appeared in Haddouche and Guedj, 2023a; Chugg et al., 2023) is that they are linked to almost surely convergence while classical PAC-Bayes results are related to in-probability convergence. We elaborate on

1. The notation $A \succeq B$ means that $A - B$ is a semi-definite positive matrix.

this fundamental difference in Appendix A.3.

To obtain time-uniform estimation bounds, we exploit the notion of flat minima, *i.e.* a minimum whose neighbourhood almost minimises the loss, and this property can be attained in an overparameterised setting such as neural networks once the optimisation phase has been performed. We exploit this flatness property through the gradient norm $\|\nabla_h \ell(h, \mathbf{z})\|$ of the loss *w.r.t.* the predictor h for any \mathbf{z} . To our knowledge, this is the first attempt to do so, as Gat et al. (2022) focus on gradients with respect to the data $\nabla_{\mathbf{z}} \ell(h, \mathbf{z})$.

We first state in Assumption 5 a key assumption of our work, which intricates the data distribution \mathcal{D} with the posterior of interest Q .

Assumption 5 *We then say that $Q \in \mathcal{M}(\mathcal{H})$ is quadratically self-bounded *w.r.t.* the loss function $\ell : \mathcal{H} \times \mathcal{Z} \rightarrow \mathbb{R}$ and the constant $C > 0$ (namely $QSB(\ell, C)$) if*

$$\mathbb{E}_{\mathbf{z} \sim \mathcal{D}} \left[\left(\mathbb{E}_{h \sim Q} [\ell(h, \mathbf{z})] \right)^2 \right] \leq CR_{\mathcal{D}}(Q) = C \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} \left[\mathbb{E}_{h \sim Q} [\ell(h, \mathbf{z})] \right].$$

Assumption 5 is a relaxation of boundedness, as if $\ell : \mathcal{H} \times \mathcal{Z} \rightarrow [0, C]$ then it is $QSB(\ell, C)$. It is an alternative to the bounded expected variance assumption in anytime-valid PAC-Bayes bounds (Haddouche and Guedj, 2023a; Chugg et al., 2023). A key issue with their boundedness assumption is that it must hold for all posteriors, including those providing poor generalisation performance. Our QSB assumption avoids this by intricately linking the properties of the distribution \mathcal{D} , the loss ℓ and the posterior Q . Such a design is in line with the conclusions of the recent work of Gastpar et al. (2024), inviting to derive generalisation bounds valid for specific pairs (Q, \mathcal{D}) (rather than uniformly valid for all such pairs) to reach sharper results. Finally, we interpret C as a contraction constant that attenuates, on average, the local expansion (governed by variances of Q and \mathcal{D}) of the loss around the mean of Q . To illustrate the applicability of QSB condition beyond bounded losses, we give a concrete example of an unbounded loss satisfying it.

Example 1 *Assume that for any $\mathbf{z} \in \mathcal{Z}$, the loss $\ell(\cdot, \mathbf{z})$ is unbounded and L -Lipschitz, that we are in the realisable case, *i.e.* there exists $h^* \in \mathcal{H}$ such that $\forall \mathbf{z} \in \mathcal{Z}$, $\ell(h^*, \mathbf{z}) = 0$, and that Q is an arbitrary distribution with mean m_Q and standard deviation σ_Q both bounded by a certain K . Then, since we have $\ell(h, \mathbf{z}) = \ell(h, \mathbf{z}) - \ell(h^*, \mathbf{z})$, and by Lipschitzness and Cauchy–Schwarz’s inequality, for any $\mathbf{z} \in \mathcal{Z}$, we can deduce that we have $\mathbb{E}_{h \sim Q} [\ell(h, \mathbf{z})]^2 \leq \left(\mathbb{E}_{h \sim Q} [\sqrt{L} \|h - h^*\| \ell(h, \mathbf{z})] \right)^2 \leq L \mathbb{E}_{h \sim Q} [\|h - h^*\|] \mathbb{E}_{h \sim Q} [\ell(h, \mathbf{z})]$. Finally, note that by Jensen’s inequality and the bias-variance decomposition, we have $\mathbb{E}_{h \sim Q} [\|h - h^*\|] \leq \sqrt{\mathbb{E}_{h \sim Q} [\|h - h^*\|^2]} = \sqrt{\sigma_Q^2 + \|m_Q - h^*\|^2} \leq \sqrt{K^2 + (K + \|h^*\|)^2}$. This ensures that the QSB condition holds in this case with constant $C = L\sqrt{K^2 + (K + \|h^*\|)^2}$.*

We are now able to state the main result of this section.

Theorem 6 *For any $C > 0$, for any λ such that $\frac{2}{C} > \lambda > 0$, for any data-free prior $P \in \mathcal{M}(\mathcal{H})$, for any loss function $\ell : \mathcal{H} \times \mathcal{Z} \rightarrow \mathbb{R}_+$, and for any $\delta \in (0, 1]$, we have, with probability at least $1 - \delta$ over*

the sample \mathcal{S} , for all $m \in \mathbb{N}^*$, for all Q being $Poinc(c_P)$, $QSB(\ell, C)$, and $\ell(\cdot, \mathbf{z}) \in H^1(Q)$ for all \mathbf{z} ,

$$\begin{aligned} R_{\mathcal{D}}(Q) \leq \frac{1}{1 - \frac{\lambda C}{2}} \left(\hat{R}_{\mathcal{S}_m}(Q) + \frac{\text{KL}(Q, P) + \log(1/\delta)}{\lambda m} \right) \\ + \frac{\lambda}{2 - \lambda C} c_P(Q) \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} \left[\mathbb{E}_{h \sim Q} (\|\nabla_h \ell(h, \mathbf{z})\|^2) \right]. \end{aligned}$$

Theorem 6 provides a time-uniform estimation bound with $\alpha = \frac{2}{\lambda(2-\lambda C)}$ and with the threshold $\varepsilon(Q, \mathcal{S}_m) = \frac{1}{2-\lambda C} \left(2\hat{R}_{\mathcal{S}_m}(Q) + \lambda c_P(Q) \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} [\mathbb{E}_{h \sim Q} (\|\nabla_h \ell(h, \mathbf{z})\|^2)] \right)$ for any $m \in \mathbb{N}^*$. Achieving a small ε_m (and thus approaching a fast convergence rate) requires two conditions: $\hat{R}_{\mathcal{S}_m}(Q) \approx 0$ and expected gradients to vanish. While the first condition is often satisfied by deep neural networks, the second holds if a flat minimum has been reached through the optimisation process. Then, setting $\lambda = 1/C$ ensures a transitory fast-rate bound of $1/m$ for any $m \in \mathbb{N}^*$. Otherwise, for a fixed m , setting $\lambda = m^{-\alpha}/C$ with $\alpha \in [0; 1/2]$ allows adapting the rate with respect to the behaviour of the gradients. In the case of constant gradients, we recover a convergence rate of $1/\sqrt{m}$, at the cost of the time-uniform property, matching [Alquier et al. \(2016, Theorem 4.1\)](#).

On the role of flat minima in PAC-Bayes learning. We highlight that the gradient term in Theorem 6 is derived with respect to a predictor $h \in \mathcal{H}$ and not $\mathbf{z} \in \mathcal{Z}$ which is, to our knowledge, novel in PAC-Bayes. This is particularly impacting, as $\nabla_h \ell(h, \mathbf{z})$ is the gradient involved in learning procedures and, when averaged over Q , provides information about the nature of the minima reached and its neighbourhood; we elaborate further in [Appendix A.4](#). Theorem 6 suggests that, to attain good generalisation ability, the mean of Q must be close to two minima: (i) on $\hat{R}_{\mathcal{S}_m}(Q)$ in order to make $\hat{R}_{\mathcal{S}_m}(Q)$ small, and (ii) on $\mathbb{E}_{\mathbf{z} \sim \mathcal{D}} [\|\nabla_h \ell(h, \mathbf{z})\|^2]$ to ensure the gradients are small. The variance of Q must fit the flatness of these minima to reduce the expected terms on the right-hand side of Theorem 6. Finally, the KL term invites, for Gaussian distributions, considering high variances and flat minima to maintain a small value for the bound.

A focus on C . Taking $\lambda = 1/C$ in Theorem 6 reduces the influence of the prior distribution P while amplifying the gradient term. Therefore, a small C is desirable when working with flat minima to mitigate the effects of a poorly chosen prior. Having a small C is reachable in practice: we show in [Section 6](#), for a classification task on MNIST, that the QSB assumption holds with C strictly smaller than 1 when considering neural networks.

Proof of Theorem 6 We start from [Chugg et al. \(2023, Corollary 17\)](#) instantiated with a single λ , an *i.i.d.* dataset and a prior P . With probability at least $1 - \delta$, for all $Q \in \mathcal{M}(\mathcal{H})$ and $m \in \mathbb{N}^*$, we have

$$R_{\mathcal{D}}(Q) \leq \hat{R}_{\mathcal{S}_m}(Q) + \frac{\text{KL}(Q, P) + \log(1/\delta)}{\lambda m} + \frac{\lambda}{2} \left(\mathbb{E}_{h \sim Q} \left[\mathbb{E}_{\mathbf{z} \sim \mathcal{D}} [\ell(h, \mathbf{z})^2] \right] \right),$$

where $\mathbf{z} \sim \mathcal{D}$ is independent from \mathcal{S} . We study the term $\mathbb{E}_{h \sim Q} [\mathbb{E}_{\mathbf{z} \sim \mathcal{D}} [\ell(h, \mathbf{z})^2]]$ on the right-hand side. We first apply Fubini's theorem to obtain

$$\mathbb{E}_{h \sim Q} \left[\mathbb{E}_{\mathbf{z} \sim \mathcal{D}} [\ell(h, \mathbf{z})^2] \right] = \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} \left[\mathbb{E}_{h \sim Q} [\ell(h, \mathbf{z})^2] \right] = \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} \left[\text{Var}_{h \sim Q} (\ell(h, \mathbf{z})) + \left(\mathbb{E}_{h \sim Q} [\ell(h, \mathbf{z})] \right)^2 \right].$$

As for any $\mathbf{z} \in \mathcal{Z}$, we have $\ell(\cdot, \mathbf{z}) \in H^1(Q)$, we apply Poincaré inequality to obtain

$$\mathbb{E}_{h \sim Q} \left[\mathbb{E}_{\mathbf{z} \sim \mathcal{D}} [\ell(h, \mathbf{z})^2] \right] \leq \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} \left[c_P(Q) \mathbb{E}_{h \sim Q} (\|\nabla_h \ell(h, \mathbf{z})\|^2) + \left(\mathbb{E}_{h \sim Q} [\ell(h, \mathbf{z})] \right)^2 \right].$$

Using that Q is $\text{QSB}(\ell, C)$ and re-organising the terms gives

$$\begin{aligned} R_{\mathcal{D}}(Q) \leq \frac{1}{1 - \frac{\lambda C}{2}} \left(\hat{R}_{\mathcal{S}_m}(Q) + \frac{\text{KL}(Q, P) + \log(1/\delta)}{\lambda m} \right) \\ + \frac{\lambda}{2 - \lambda C} c_P(Q) \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} \left[\mathbb{E}_{h \sim Q} (\|\nabla_h \ell(h, \mathbf{z})\|^2) \right]. \end{aligned}$$

■

When the QSB assumption is not verified, it is still possible to exploit the benefit of flat minima in PAC-Bayes at the cost of an upper bound on $R_{\mathcal{D}}(Q)$ and a supplementary Poincaré assumption on the data distribution \mathcal{D} .

Corollary 7 *For any $C > 0$, for any λ such that $\frac{2}{C} > \lambda > 0$, for any data-free prior $P \in \mathcal{M}(\mathcal{H})$, for any loss function $\ell : \mathcal{H} \times \mathcal{Z} \rightarrow \mathbb{R}_+$ such that $\ell(h, \cdot)$ is C^1 almost everywhere on \mathcal{Z} , any $\delta \in (0, 1]$, if the data distribution \mathcal{D} is $\text{Poinc}(c_P)$, with probability at least $1 - \delta$ over the sample \mathcal{S} , for any $m \in \mathbb{N}^*$, any posterior Q being $\text{Poinc}(c_P)$ with $R_{\mathcal{D}}(Q) \leq C$ and such that for any $\mathbf{z} \in \mathcal{Z}$, $\ell(\cdot, \mathbf{z}) \in H^1(Q)$:*

$$\begin{aligned} R_{\mathcal{D}}(Q) \leq \frac{1}{1 - \frac{\lambda C}{2}} \left(\hat{R}_{\mathcal{S}_m}(Q) + \frac{\text{KL}(Q, P) + \log(1/\delta)}{\lambda m} \right) \\ + \frac{\lambda}{2 - \lambda C} \left(c_P(Q) \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} \left[\mathbb{E}_{h \sim Q} (\|\nabla_h \ell(h, \mathbf{z})\|^2) \right] + c_P(\mathcal{D}) \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} \left(\left\| \mathbb{E}_{h \sim Q} [\nabla_{\mathbf{z}} \ell(h, \mathbf{z})] \right\|^2 \right) \right). \end{aligned}$$

The proof is deferred to Appendix C.1. Corollary 7 states that if Q reaches a flat minimum (meaning $\|\nabla_h \ell(h, \mathbf{z})\|$ is small) and this minimum is robust to the training dataset (meaning $\|\nabla_{\mathbf{z}} \ell(h, \mathbf{z})\|$ is small), then a time-uniform estimation bound is attainable with small ε_m , requiring only an upper bound on $R_{\mathcal{D}}(Q)$. The assumption of small $\|\nabla_{\mathbf{z}} \ell(h, \mathbf{z})\|$ can be attained with algorithms such as Sharpness-Aware Minimisation (Foret et al., 2021). Another lead would be to focus on specific predictors where this gradient is directly controlled, as in Lipschitz neural networks training or in adversarial robustness training (Madry et al., 2018; Li et al., 2019).

Corollary 7 holds when \mathcal{D} is $\text{Poinc}(c_P)$, encompassing the case of Gaussian mixtures (Schlichting, 2019), which can approximate any smooth density (as recalled in Gat et al., 2022). However, the Poincaré constant of a general mixture is unknown, and the upper bound of Schlichting (2019) scales with the number of components, involving potentially high χ^2 divergences.

Comparison with Gat et al. (2022). We compare Corollary 7 with Gat et al. (2022, Theorems 3.5 and 3.6). First, our result holds under the assumption that the distribution \mathcal{D} follows a Poincaré inequality, which is strictly less restrictive than assuming a log-Sobolev inequality (Proposition 15). Second, they assume a bounded loss and focus solely on classification tasks satisfying a technical

assumption (see their Lemma 3.3) while ours holds for any learning problem at the sole assumption of a bounded $R_{\mathcal{D}}(Q)$, allowing the loss ℓ to be unbounded (and non-negative). To conclude their proof, Gat et al. (2022) use a uniform bound on $\mathbb{E}_{\mathbf{z} \sim \mathcal{D}}[\|\nabla_{\mathbf{z}}\ell(h, \mathbf{z})\|]$ in their Theorem 3.5 to have a tractable bound, diminishing the benefits of gradient norm. While they address this limitation in Gat et al. (2022, Theorem 3.6), the explicit influence of the gradient norm appears within an exponential moment on the losses (attenuated by a logarithm), averaged *w.r.t.* the data-free prior P . Thus, the associated gradients have no apparent reason to be small, and their result cannot be linked to flat minima. In contrast, Corollary 7 involves expected gradients *w.r.t.* the posterior Q , which may reach flat minima.

3.2. Towards fully empirical bound for gradient-Lipschitz functions

In this section, we assume that the gradient $\nabla_h \ell(h, \mathbf{z})$ is G -Lipschitz for any $\mathbf{z} \in \mathcal{Z}$, which is a classical assumption in optimisation, especially when considering non-convex objectives for SGD (Ghadimi and Lan, 2013; Panageas and Piliouras, 2017; Garrigos and Gower, 2023). A large portion of high-probability PAC-Bayes bounds are fully empirical, meaning that the right-hand side of the bounds can be computed. This has numerous advantages, including in-training numerical evaluation of the bound and the development of novel PAC-Bayesian algorithms that minimises such empirical bounds; see (Dziugaite and Roy, 2017; Pérez-Ortiz et al., 2021; Viallard et al., 2023b) among others. However, Theorem 6 and Corollary 7 are not fully empirical, as they involve terms such as $\mathbb{E}_{\mathbf{z} \sim \mathcal{D}}[\mathbb{E}_{h \sim Q}[\|\nabla_h \ell(h, \mathbf{z})\|^2]]$ and $\mathbb{E}_{\mathbf{z} \sim \mathcal{D}}[\|\mathbb{E}_{h \sim Q}[\nabla_{\mathbf{z}} \ell(h, \mathbf{z})]\|^2]$, which involves an expectation over $\mathbf{z} \sim \mathcal{D}$ and are thus not computable in practice. As a result, they lack the desirable properties of fully empirical bounds; we address this issue in Theorem 8.

Theorem 8 *For any $C_1, C_2, c > 0$, for any data-free prior $P \in \mathcal{M}(\mathcal{H})$, for any loss function $\ell : \mathcal{H} \times \mathcal{Z} \rightarrow \mathbb{R}_+$ being C^2 and for any $\delta \in (0, 1]$, we have, with probability at least $1 - \delta$ over the sample \mathcal{S} , for all $m \in \mathbb{N}^*$, for all Q being $\text{Poinc}(c_P)=c$, $\text{QSB}(\ell, C_1)$, $\text{QSB}(\|\nabla_h \ell\|^2, C_2)$, and $\ell(\cdot, \mathbf{z}) \in H^1(Q)$, and $\|\nabla_h \ell(\cdot, \mathbf{z})\|^2 \in H^1(Q)$ for all \mathbf{z} ,*

$$R_{\mathcal{D}}(Q) \leq 2\hat{R}_{\mathcal{S}_m}(Q) + \frac{2c}{C_1} \mathbb{E}_{h \sim Q} \left[\frac{1}{m} \sum_{i=1}^m \|\nabla_h \ell(h, \mathbf{z}_i)\|^2 \right] + 2 \left(C_1 + c \frac{4cG^2 + C_2}{C_1} \right) \frac{\text{KL}(Q, P) + \log(2/\delta)}{m}.$$

The proof is deferred to Appendix C.2. We showed that to attain fast rates, the QSB assumption must hold for both the loss and its gradient. We are then able to derive an empirical generalisation bound, involving both empirical loss and gradients.

As Theorem 8 is fully empirical, it can be transformed into a generalisation metric, *i.e.* an empirical function of the predictor whose increase or decrease is correlated to the increase or decrease of the generalisation ability of the predictor. In the case of PAC-Bayes, the generalisation metric comes from the generalisation bound. Such an idea has been exploited recently (Neyshabur et al., 2017; Jiang et al., 2020; Dziugaite et al., 2020; Viallard et al., 2024a) to show that flatness of the empirical risk was correlated to generalisation. In particular, from $\hat{R}_{\mathcal{S}}(Q)$, Neyshabur et al. (2017) derived a notion of *sharpness*, stated in Equation (1), which gives information about the flatness of

the reached minima for any $Q = \mathcal{N}(\mu_Q, \sigma^2 \text{Id})$. This notion is defined by

$$\mathbb{E}_{\nu \sim \mathcal{N}(\mathbf{0}, \sigma^2 \text{Id})} \left[\hat{R}_{\mathcal{S}_m}(\mu_Q + \nu) - \hat{R}_{\mathcal{S}_m}(\mu_Q) \right]. \quad (1)$$

Sharpness is then the averaged risk of a predictor drawn under $\mathcal{N}(\mu_Q, \sigma^2 \text{Id})$ w.r.t. its mean Q . Theorem 8 enhance this notion of sharpness with empirical gradients when Q is $QSB(\ell, C_1)$:

$$\text{Sharp}_{\frac{\sigma^2}{C_1}}(Q) := \mathbb{E}_{\nu \sim \mathcal{N}(\mathbf{0}, \sigma^2 \text{Id})} \left[2 \left[\hat{R}_{\mathcal{S}_m}(\mu_Q + \nu) - \hat{R}_{\mathcal{S}_m}(\mu_Q) \right] + \frac{\sigma^2}{C_1} \left[\hat{G}_{\mathcal{S}_m}(\mu_Q + \nu) - \hat{G}_{\mathcal{S}_m}(\mu_Q) \right] \right], \quad (2)$$

where $\hat{G}_{\mathcal{S}_m}(h) = \frac{1}{m} \sum_{i=1}^m \|\nabla_h \ell(h, \mathbf{z}_i)\|^2$. This gradient term can be seen as the norm of an empirical Fisher information, linked to the second-order moment derivative. Thus, Equation (2) involves a notion of flatness on both the loss and its gradient, unlike Equation (1). For the sake of clarity, we specialise Theorem 8 in Corollary 9 to Gaussian distributions, introducing this notion of sharpness.

Corollary 9 *For any $C_1, C_2 > 0$, for any data-free prior $P = \mathcal{N}(\mu_P, \sigma^2 \text{Id})$ with fixed variance $\sigma^2 > 0$, for any loss function $\ell : \mathcal{H} \times \mathcal{Z} \rightarrow \mathbb{R}_+$ being C^2 for and any $\delta \in (0, 1]$, we have, with probability at least $1 - \delta$ over the sample \mathcal{S} , for all $m \in \mathbb{N}^*$, for all $Q = \mathcal{N}(\mu_Q, \sigma^2 \text{Id})$ being $QSB(\ell, C_1)$, $QSB(\|\nabla_h \ell\|^2, C_2)$ and $\ell(\cdot, \mathbf{z}) \in H^1(Q)$, and $\|\nabla_h \ell(\cdot, \mathbf{z})\|^2 \in H^1(Q)$ for all \mathbf{z} ,*

$$R_{\mathcal{D}}(Q) \leq 2\hat{R}_{\mathcal{S}_m}(\mu_Q) + \hat{G}_{\mathcal{S}_m}(\mu_Q) + \text{Sharp}_{\frac{\sigma^2}{C_1}}(Q) + \mathcal{O}\left(\frac{\text{KL}(Q, P) + \log(2/\delta)}{m}\right).$$

4. Generalisation ability of Gibbs distributions with a log-Sobolev prior

A limitation of the results in Section 3 is that the KL divergence term remains generally uncontrolled, as its formulation depends on the nature of P and Q . While a closed form exists for Gaussian distributions, a natural question is whether it is possible to explicitly control the KL term for another class of distributions. Following the approach of Catoni (2007), we focus in this section on Gibbs posteriors, which naturally arise in PAC-Bayes through the use of tools from statistical physics. We show that log-Sobolev inequalities allow us to control the KL divergence of such distributions with respect to their priors.

Controlling the KL divergence when Q is a Gibbs posterior. Lemma 10 exploits the fact that the KL divergence can be expressed as an entropy with respect to the prior distribution P . It shows that the KL divergence of the Gibbs posterior $Q_{\mathcal{S}_m}^\gamma$ is upper-bounded by gradient terms, provided that the prior P satisfies a log-Sobolev inequality.

Lemma 10 *For any $\gamma > 0$, for any $m \in \mathbb{N}^*$, for any data-free prior $P \in \mathcal{M}(\mathcal{H})$ being L -Sob(c_{LS}), for any loss function $\ell : \mathcal{H} \times \mathcal{Z} \rightarrow \mathbb{R}_+$ such that $\ell(\cdot, \mathbf{z}) \in H^1(P)$ for any \mathbf{z} , we have*

$$\text{KL}(Q_{\mathcal{S}_m}^\gamma, P) \leq \frac{\gamma^2 c_{LS}(P)}{4} \mathbb{E}_{h \sim Q_{\mathcal{S}_m}^\gamma} \left[\|\nabla_h \hat{R}_{\mathcal{S}_m}(h)\|^2 \right].$$

The proof is deferred to Appendix C.3. The key message of this lemma is that for Gibbs posteriors, the expansion of the KL divergence is controlled by an expected empirical gradient term. Note in this case that, while the KL divergence has an explicit formulation, it requires calculating the exponential moment $\mathbb{E}_{h \sim P}[\exp(-\gamma \hat{R}_{\mathcal{S}_m}(h))]$ which is costly in practice. In contrast, we only need to estimate a second-order moment over $Q_{\mathcal{S}_m}^\gamma$.

Generalisation ability of Gibbs posteriors. When Gibbs posteriors are involved, the KL divergence can be controlled by a gradient term. An ideal approach, as in Section 3, would be to involve a Poincaré inequality. However, Gibbs posteriors do not necessarily satisfy a Poincaré inequality as in Section 3, so we need to make supplementary assumptions about the loss.

Theorem 11 *We first have $C > 0$, $\gamma > 0$, and a data-free prior $P \in \mathcal{M}(\mathcal{H})$ being L -Sob(c_{LS}). (i) For any loss function $\ell : \mathcal{H} \times \mathcal{Z} \rightarrow [0, 1]$ and for any $\delta \in (0, 1]$, with probability at least $1 - \delta$ over the sample \mathcal{S} , for any $m \in \mathbb{N}^*$, we have*

$$R_{\mathcal{D}}(Q_{\mathcal{S}_m}^{\gamma}) \leq 2 \left(\hat{R}_{\mathcal{S}_m}(Q_{\mathcal{S}_m}^{\gamma}) + \frac{\gamma^2 c_{LS}(P)}{4m} \mathbb{E}_{h \sim Q_{\mathcal{S}_m}^{\gamma}} \left[\|\nabla_h \hat{R}_{\mathcal{S}_m}(h)\|^2 \right] + \frac{\log(1/\delta)}{m} \right).$$

(ii) *For any loss function ℓ and prior P satisfying the conditions of Proposition 4. Then, for any $\frac{2}{C} > \lambda > 0$, with probability at least $1 - \delta$ over the sample \mathcal{S} , for all $m \in \mathbb{N}^*$, assuming $\ell(\cdot, \mathbf{z}) \in \mathbb{H}^1(Q_{\mathcal{S}_m}^{\gamma})$ and that $Q_{\mathcal{S}_m}^{\gamma}$ is QSB(C), we have*

$$R_{\mathcal{D}}(Q_{\mathcal{S}_m}^{\gamma}) \leq \frac{1}{1 - \frac{\lambda C}{2}} \left(\hat{R}_{\mathcal{S}_m}(Q_{\mathcal{S}_m}^{\gamma}) + \frac{\gamma^2 c_{LS}(P)}{4\lambda m} \mathbb{E}_{h \sim Q_{\mathcal{S}_m}^{\gamma}} \left[\|\nabla_h \hat{R}_{\mathcal{S}_m}(h)\|^2 \right] + \frac{\log(1/\delta)}{\lambda m} \right) + \frac{\lambda e^{4\|\ell_2\|_{\infty}} c_{LS}(P)}{4 - 2\lambda C} \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} \left[\mathbb{E}_{h \sim Q_{\mathcal{S}_m}^{\gamma}} (\|\nabla_h \ell(h, \mathbf{z})\|^2) \right].$$

The proof is deferred to Appendix C.4. Note that we could have also derived a result analogous to Corollary 7 at the cost of an additional Poincaré assumption on \mathcal{D} . The influence of the inverse temperature γ is quadratic: this is the price to pay to fit the dataset and reduce the influence of the prior. This dependency is attenuated by a gradient term, which is small if a flat minimum on the $\hat{R}_{\mathcal{S}_m}(h)$ has been reached. This suggests that in the case of Gibbs posteriors with a log-Sobolev prior, reaching a flat minimum on $\hat{R}_{\mathcal{S}_m}(h)$ controls not only $\hat{R}_{\mathcal{S}_m}(Q)$, but also the KL divergence and this last property is not reachable when considering Poincaré distributions. The other gradient term comes from Section 3 and requires to be close to a flat minimum on $R_{\mathcal{D}}(h)$ to attain fast rates.

Comparison to literature. To our knowledge, existing bounds for Gibbs posteriors do not involve the gradient norm of the posterior. For instance, Zhang (2006, Theorem 4.2) involves an exponential moment with respect to the prior distribution, while Kuzborskij et al. (2019, Equation 8) consider flatness through ellipsoids around a minimum, controlling it via an 'effective dimension', which is a function of the eigenvalue of the Hessian of the theoretical risk. Finally, Rivasplata et al. (2020, Equation 6) propose a bound for bounded losses, removing the KL divergence when Gibbs posteriors are considered, with a rate of $\mathcal{O}(1/\sqrt{m} + \gamma/m)$. Assuming the Gibbs posterior reaches a flat minimum such that $\gamma \mathbb{E}_{h \sim Q_{\mathcal{S}_m}^{\gamma}} [\|\nabla_h \hat{R}_{\mathcal{S}_m}(h)\|^2] \leq 1$, then Theorem 11 yields a bound of magnitude $\mathcal{O}(\hat{R}_{\mathcal{S}_m}(Q_{\mathcal{S}_m}^{\gamma}) + 1/m + \gamma/m)$, representing a transitory fast rate with threshold $\hat{R}_{\mathcal{S}_m}(Q_{\mathcal{S}_m}^{\gamma})$.

5. On the benefits of the gradient norm in Wasserstein PAC-Bayes learning

In Sections 3 and 4, we provided various generalisation bounds, benefiting from flat minima. However, our results involve a KL divergence, implying absolute continuity of Q with respect to P , making them incompatible with deterministic predictors (obtained with Dirac distributions). Then

the following question arises: can we benefit from the desirable properties of flat minima in a PAC-Bayes bound valid for deterministic predictors? To address this issue, a recent line of work has emerged, focusing on integral probability metrics, with particular attention to the 1-Wasserstein distance (Amit et al., 2022; Haddouche and Guedj, 2023b; Viillard et al., 2023b, 2024b); see Definition 16. The idea behind these works is to replace the change of measure inequality (Csiszár, 1975; Donsker and Varadhan, 1976) with the Kantorovich-Rubinstein duality (Villani, 2009), trading a KL divergence for a Wasserstein distance. This not only yields sound theoretical bounds but also PAC-Bayes algorithms for deterministic predictors and Lipschitz neural networks (Viillard et al., 2023b). We go further here by obtaining the first PAC-Bayesian bound directly involving a 2-Wasserstein distance, trading a Lipschitz assumption for a gradient-Lipschitz one, which is well-suited for optimisation. To do so, we derive a novel change of measure inequality tailored for the condition (\star) described below, which is a relaxation of the gradient Lipschitz assumption.

$f : \mathcal{H} \rightarrow \mathbb{R}$ satisfies $(\star) \Leftrightarrow \exists G > 0, \forall (a, a') \in \mathcal{H}^2, \langle \nabla f(\cdot), a - a' \rangle$ is $G\|a - a'\|$ -Lipschitz.

Theorem 12 *For any predictor set \mathcal{H} with finite diameter $D > 0$, and for any function $f : \mathcal{H} \rightarrow \mathbb{R}$ satisfying (\star) , we have for all distributions $P \in \mathcal{M}(\mathcal{H})$ and $Q \in \mathcal{M}(\mathcal{H})$*

$$\mathbb{E}_{h \sim Q} [f(h)] \leq \frac{G}{2} W_2^2(Q, P) + \mathbb{E}_{h \sim P} [f(h)] + D \mathbb{E}_{h \sim Q} [\|\nabla f(h)\|].$$

The proof is deferred to Appendix C.5. Theorem 12 shows that, when gradients are Lipschitz, it is possible to obtain a duality formula involving the gradient of the considered function, at the cost of a linear dependency on the diameter D of \mathcal{H} . Theorem 12 is also linked to the change of measure inequality (Csiszár, 1975; Donsker and Varadhan, 1976) when the prior distribution satisfies a log-Sobolev inequality. This link is detailed in Corollary 13.

Corollary 13 *For any distribution P being L -Sob(c_{LS}) such that $dP(h) \propto \exp(-V(h))dh$, with V being \mathcal{C}^2 , for any $R > 0$, for any function f on the centred ball $\mathcal{B}(\mathbf{0}, R)$ of radius R satisfying the (\star) assumption, and for any distribution $Q \in \mathcal{M}(\mathcal{H})$, we have:*

$$\mathbb{E}_{h \sim Q} [f(\mathcal{P}_R(h))] \leq \frac{G c_{LS}(P)}{4} \text{KL}(Q, P) + \mathbb{E}_{h \sim P} [f(\mathcal{P}_R(h))] + 2R \mathbb{E}_{h \sim Q} [\|\nabla_h f(\mathcal{P}_R(h))\|],$$

where \mathcal{P}_R denotes the Euclidean projection onto $\mathcal{B}(\mathbf{0}, R)$.

The proof is deferred to Appendix C.6. Corollary 13 involves a KL divergence and an Euclidean predictor space $\mathcal{H} = \mathbb{R}^d$. This comes at the cost of approximating Q and P by, respectively, $\mathcal{P}_R\#Q$ and $\mathcal{P}_R\#P$. Thus, the radius R is now a hyperparameter balancing the tradeoff between the quality of our approximations and the looseness of the bound (if the gradient norm is large). A notable strength is that the smoothness assumption is relaxed to apply only within the centred ball $\mathcal{B}(\mathbf{0}, R)$.

From Theorem 12, we now derive a novel generalisation bound allowing deterministic predictors.

Theorem 14 *Let $\delta \in (0, 1)$ and $P \in \mathcal{M}(\mathcal{H})$ a data-free prior. Assume \mathcal{H} has a finite diameter $D > 0$, for any loss function $\ell : \mathcal{H} \times \mathcal{Z} \rightarrow \mathbb{R}_+$ and any $m \in \mathbb{N}^*$, the generalisation gap $h \mapsto$*

$R_{\mathcal{D}}(h) - \hat{R}_{\mathcal{S}_m}(h)$ satisfies (\star) . Assume that $\mathbb{E}_{h \sim P} \mathbb{E}_{\mathbf{z} \sim \mathcal{D}}[\ell(h, \mathbf{z})^2] \leq \sigma^2$, then the following holds with probability at least $1 - \delta$, for any $m > 0$ and any $Q \in \mathcal{M}(\mathcal{H})$:

$$R_{\mathcal{D}}(Q) \leq \hat{R}_{\mathcal{S}_m}(Q) + \frac{G}{2} W_2^2(Q, P) + \sqrt{\frac{2\sigma^2 \log(\frac{1}{\delta})}{m}} + D \mathbb{E}_{h \sim Q} \left(\left\| \nabla_h R_{\mathcal{D}}(h) - \nabla_h \hat{R}_{\mathcal{S}_m}(h) \right\| \right).$$

The proof is deferred to Appendix C.7. Theorem 14 is not the first generalisation bound to involve a 2-Wasserstein distance (Lugosi and Neu, 2022, 2023). However, these results require infinitely smooth loss functions. Additionally, the results from Amit et al. (2022); Haddouche and Guedj (2023b); Viallard et al. (2023b), which use the 1-Wasserstein, can be directly relaxed on bounds involving the 2-Wasserstein, while still requiring a Lipschitz loss. In contrast, our result holds for any nonnegative loss whom the generalisation gap $h \rightarrow R_{\mathcal{D}}(h) - \hat{R}_{\mathcal{S}_m}(h)$, satisfies (\star) , being a relaxation of gradient-Lipschitz assumption. Theorem 14 involves a rate of $1/\sqrt{m}$, as we have to control the generalisation gap over P . Another restriction of our result, compared to previous ones, is that it holds for \mathcal{H} with a finite diameter.

Can Theorem 14 go to zero with large m ? In its current form, it is unclear whether Theorem 14 goes to zero with large m , as the Wasserstein distance and the gradient term, have no explicit convergence rate in m . Concerning the gradient term, it has been shown that, if we consider Q to be a Dirac in the output of SGD after T iterations, then according to Li and Liu (2022, Theorem 3), when the intrinsic noise of SGD is subgaussian, the output $\mathbf{w}_T \in \mathbb{R}^d$ of SGD satisfies with high probability: $\|\nabla_h R_{\mathcal{S}_m}(\mathbf{w}_T) - \nabla_h R_{\mathcal{D}}(\mathbf{w}_T)\|^2 \leq \mathcal{O}\left(\frac{d\sqrt{T}\log(T)}{m}\right)$. Concerning the Wasserstein term, if we assume directly that ℓ satisfies (\star) with constant G , then using a technique inspired from Amit et al. (2022) yields with probability $1 - \delta$ that the generalisation gap satisfies (\star) with constant $G' = \mathcal{O}\left(G\sqrt{\log(|\mathcal{H}|/\delta)/m}\right)$ when \mathcal{H} is finite. We prove this in Appendix A.5.

6. An empirical study of Assumption 5 for neural networks²

In this section, we empirically verify whether the QSB assumption holds for neural nets. This allows us to assess whether Theorem 6 helps in understanding the generalisation ability of neural nets.

Experimental protocol. We consider classification tasks on two datasets: MNIST (LeCun, 1998) and FashionMNIST (Xiao et al., 2017). We have kept the original training set \mathcal{S}_m and the original test set denoted by \mathcal{T}_n (of size n). We consider the convolutional neural network of Springenberg et al. (2015) adapted for MNIST and FashionMNIST. The model is composed of 4 layers containing 10 channels with a 5×5 -kernel; we set the stride and the padding to 1, except for the second layer, where it is fixed to 2. Each of these (convolutional) layers is followed by a Leaky ReLU activation function. Moreover, an average pooling with a 8×8 -kernel is performed before the Softmax activation function. To initialise the weights of the network, we use Glorot and Bengio (2010) uniform initialiser, while the biases are initialised in $[-1/\sqrt{250}, +1/\sqrt{250}]$ uniformly (except the first layer, the interval is $[-1/5, +1/5]$). Hence, in this case, \mathcal{H} is the set of neural networks with a fixed architecture, and parametrised with a vector \mathbf{w} . The posterior distribution Q is a Gaussian measure $\mathcal{N}(\mathbf{w}, \sigma^2 \text{Id})$ centred on the parameters \mathbf{w} associated with the model; σ is set to 10^{-4} . Note that this distribution

2. The source code is available at this [link](#).

respects the $\text{Poinc}(c_P)$ assumption; see Section 3.1. We train the neural network with the (vanilla) stochastic gradient descent algorithm, where the batch size is equal to 512, and the learning rate is fixed to 10^{-2} . We train for at least 10^4 gradient steps and finish the current epoch when this number of iterations is reached. Our loss ℓ is the bounded cross-entropy loss of Dziugaite and Roy (2017, Section D).

In Figure 1, we report the evolution of three quantities: (i) the estimated value of C , (ii) the test risk $\hat{R}_{\mathcal{T}_n}(\mathbb{Q})$ and (iii) the test risk with the 01-loss. More precisely, the risks and C are estimated by sampling 10 hypotheses from \mathbb{Q} and by computing the values on a mini-batch of \mathcal{T}_n (with 512 examples) at each iteration. Then, Figure 1 represents averaged values on 5 runs, each point of the curve representing the average on 100 iterations of the training process (for 10^4 iterations, we only plot 10^2 averaged points for clarity).

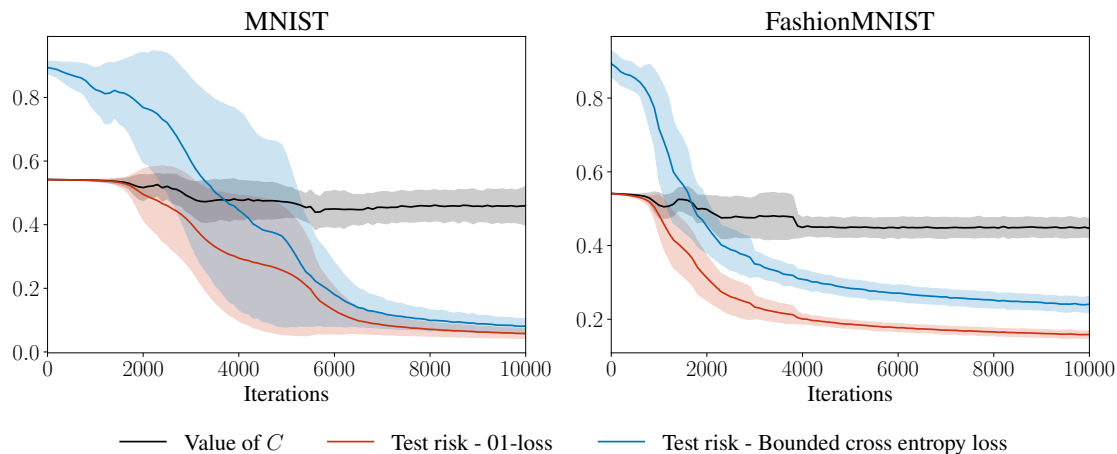


Figure 1: Evolution of the test risks (with the 01-loss and the bounded cross-entropy loss) and the value of C during the training phase.

Empirical findings. Figure 1 illustrates that, when neural networks are involved for two classification tasks, \mathbb{Q} evolves during the optimisation process while maintaining the \mathbb{Q}_{SB} property with constant $C < 1$. For both MNIST and FashionMNIST, the constant C decreases from approximately 0.55 to 0.45. We deduce that having a data-free \mathbb{P} (0 iteration) being \mathbb{Q}_{SB} with $C < 1$ suggests that the architecture of our neural network also has an influence on the \mathbb{Q}_{SB} assumption. As specified in Section 3, having $C < 1$ attenuates the impact of the KL term, thus \mathbb{P} . This is desirable as it allows the optimiser to deeply explore the predictor space when \mathbb{P} yields poor performances. We also note that the generalisation ability of \mathbb{Q} on the training loss nearly matches the performance on the 0-1 loss for MNIST but is deteriorated for FashionMNIST, this invites to study more deeply the design of such surrogates in future work.

Finally, the take-home message of this study is that the \mathbb{Q}_{SB} assumption is verified for small neural networks on MNIST. Such an empirical confirmation is crucial as it is required for our main

result (Theorem 6) and thus confirms that, for neural networks, reaching flat minima during the optimisation phase translates in increased generalisation ability.

7. Conclusion

We provide novel time-uniform PAC-Bayes bounds, that can be interpreted as a transitory convergence rate of $1/m$ when a low empirical error is reached and expected gradients vanish. Doing so, we draw sound theoretical links highlighting the impact of flat minima in generalisation. However, a crucial open question remains: how do optimisation algorithms successfully attain flat minima in the overparameterised setting? We leave this important question for future work.

Acknowledgments

M. Haddouche and U. Şimşekli are supported by the European Union (ERC grant DYNASTY 101039676). The French government partly funded this work under the management of Agence Nationale de la Recherche as part of the “France 2030” program, reference ANR-23-IACL-0008 (PR[AI]RIE-PSAI). B. Guedj acknowledges partial support from the French National Agency for Research, through the programme “France 2030” and PEPR IA on grant SHARP ANR-23-PEIA-0008. We would like to thank all the reviewers for the fruitful discussion phase, which greatly enhanced this work.

References

- Pierre Alquier. User-friendly Introduction to PAC-Bayes Bounds. *Foundations and Trends® in Machine Learning*, 2024.
- Pierre Alquier, James Ridgway, and Nicolas Chopin. On the properties of variational approximations of Gibbs posteriors. *Journal of Machine Learning Research*, 2016.
- Ron Amit, Baruch Epstein, Shay Moran, and Ron Meir. Integral Probability Metrics PAC-Bayes Bounds. In *Conference on Neural Information Processing Systems*, 2022.
- Maksym Andriushchenko, Francesco Croce, Maximilian Müller, Matthias Hein, and Nicolas Flammarion. A Modern Look at the Relationship between Sharpness and Generalization. In *International Conference on Machine Learning*, 2023.
- Cécile Ané, Sébastien Blachère, Djalil Chafaï, Pierre Fougères, Ivan Gentil, Florent Malrieu, Cyril Roberto, and Grégory Scheffer. *Sur les inégalités de Sobolev logarithmiques*. Société mathématique de France, 2000.
- William Beckner. A Generalized Poincaré Inequality for Gaussian Measures. *Proceedings of the American Mathematical Society*, 1989.
- Peter Bickel and Kjell Doksum. *Mathematical statistics: basic ideas and selected topics*. CRC Press, 2015.
- Herm Brascamp and Elliott Lieb. On extensions of the Brunn-Minkowski and Prékopa-Leindler theorems, including inequalities for log concave functions, and with an application to the diffusion equation. *Journal of functional analysis*, 1976.

- Olivier Catoni. *PAC-Bayesian supervised classification: the thermodynamics of statistical learning*. Institute of Mathematical Statistics, 2007.
- Djalil Chafaï. Entropies, convexity, and functional inequalities, On Φ -entropies and Φ -Sobolev inequalities. *Journal of Mathematics of Kyoto University*, 2004.
- Ben Chugg, Hongjian Wang, and Aaditya Ramdas. A Unified Recipe for Deriving (Time-Uniform) PAC-Bayes Bounds. *Journal of Machine Learning Research*, 2023.
- I. Csiszár. *I-Divergence Geometry of Probability Distributions and Minimization Problems*. *The Annals of Probability*, 1975.
- M. D. Donsker and S. R. S. Varadhan. Asymptotic evaluation of certain Markov process expectations for large time—III. *Communications on Pure and Applied Mathematics*, 1976.
- Gintare Dziugaite and Daniel Roy. Computing Nonvacuous Generalization Bounds for Deep (Stochastic) Neural Networks with Many More Parameters than Training Data. In *Conference on Uncertainty in Artificial Intelligence*, 2017.
- Gintare Dziugaite, Alexandre Drouin, Brady Neal, Nitarshan Rajkumar, Ethan Caballero, Linbo Wang, Ioannis Mitliagkas, and Daniel Roy. In search of robust measures of generalization. In *Conference on Neural Information Processing Systems*, 2020.
- Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. Sharpness-aware Minimization for Efficiently Improving Generalization. In *International Conference on Learning Representations*, 2021.
- Guillaume Garrigos and Robert M. Gower. Handbook of Convergence Theorems for (Stochastic) Gradient Methods. *arXiv*, abs/2301.11235, 2023.
- Michael Gastpar, Ido Nachum, Jonathan Shafer, and Thomas Weinberger. Fantastic Generalization Measures are Nowhere to be Found. In *International Conference on Learning Representations*, 2024.
- Itai Gat, Yossi Adi, Alexander Schwing, and Tamir Hazan. On the Importance of Gradient Norm in PAC-Bayesian Bounds. In *Conference on Neural Information Processing Systems*, 2022.
- Saeed Ghadimi and Guanghui Lan. Stochastic First- and Zeroth-order Methods for Nonconvex Stochastic Programming. *SIAM Journal on Optimization*, 2013.
- Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *International Conference on Artificial Intelligence and Statistics*, 2010.
- Leonard Gross. Logarithmic Sobolev Inequalities. *American Journal of Mathematics*, 1975.
- Benjamin Guedj. A Primer on PAC-Bayesian Learning. In *Proceedings of the second congress of the French Mathematical Society*, 2019.
- Alice Guionnet and Bogusław Zegarliński. Lectures on Logarithmic Sobolev Inequalities. *Séminaire de probabilités XXXVI*, 2003.

- Maxime Haddouche and Benjamin Guedj. Online PAC-Bayes Learning. In *Conference on Neural Information Processing Systems*, 2022.
- Maxime Haddouche and Benjamin Guedj. PAC-Bayes Generalisation Bounds for Heavy-Tailed Losses through Supermartingales. *Transactions on Machine Learning Research*, 2023a.
- Maxime Haddouche and Benjamin Guedj. Wasserstein PAC-Bayes Learning: A Bridge Between Generalisation and Optimisation. *arXiv*, abs/2304.07048, 2023b.
- Fredrik Hellström, Giuseppe Durisi, Benjamin Guedj, and Maxim Raginsky. Generalization Bounds: Perspectives from Information Theory and PAC-Bayes. *arXiv*, abs/2309.04381, 2023.
- Sepp Hochreiter and Jürgen Schmidhuber. Flat minima. *Neural computation*, 1997.
- Stanislaw Jastrzebski, Zachary Kenton, Devansh Arpit, Nicolas Ballas, Asja Fischer, Yoshua Bengio, and Amos J. Storkey. Three Factors Influencing Minima in SGD. *arXiv*, abs/1711.04623, 2017.
- Yiding Jiang, Behnam Neyshabur, Hossein Mobahi, Dilip Krishnan, and Samy Bengio. Fantastic Generalization Measures and Where to Find Them. In *International Conference on Learning Representations*, 2020.
- Ilya Kuzborskij, Nicolò Cesa-Bianchi, and Csaba Szepesvári. Distribution-Dependent Analysis of Gibbs-ERM Principle. In *Conference on Learning Theory*, 2019.
- Yann LeCun. The MNIST database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998.
- Michel Ledoux. Concentration of measure and logarithmic Sobolev inequalities. In *Séminaire de Probabilités XXXIII*, 2006.
- Qiyang Li, Saminul Haque, Cem Anil, James Lucas, Roger B. Grosse, and Jörn-Henrik Jacobsen. Preventing Gradient Attenuation in Lipschitz Constrained Convolutional Networks. In *Conference on Neural Information Processing Systems*, 2019.
- Shaojie Li and Yong Liu. High Probability Guarantees for Nonconvex Stochastic Gradient Descent with Heavy Tails. In *International Conference on Machine Learning*, 2022.
- Gábor Lugosi and Gergely Neu. Generalization Bounds via Convex Analysis. In *Conference on Learning Theory*, 2022.
- Gábor Lugosi and Gergely Neu. Online-to-PAC Conversions: Generalization Bounds via Regret Analysis. *arXiv*, abs/2305.19674, 2023.
- Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards Deep Learning Models Resistant to Adversarial Attacks. In *International Conference on Learning Representations*, 2018.
- David A. McAllester. Some PAC-Bayesian Theorems. *Machine Learning*, 1999.
- Jeffrey Negrea, Mahdi Haghifam, Gintare Karolina Dziugaite, Ashish Khisti, and Daniel M. Roy. Information-Theoretic Generalization Bounds for SGLD via Data-Dependent Estimates. In *Conference on Neural Information Processing Systems*, 2019.

- Gergely Neu. Information-Theoretic Generalization Bounds for Stochastic Gradient Descent. In *Conference on Learning Theory*, 2021.
- Behnam Neyshabur, Srinadh Bhojanapalli, David McAllester, and Nati Srebro. Exploring Generalization in Deep Learning. In *Conference on Neural Information Processing Systems*, 2017.
- F. Otto and C. Villani. Generalization of an Inequality by Talagrand and Links with the Logarithmic Sobolev Inequality. *Journal of Functional Analysis*, 2000.
- Ioannis Panageas and Georgios Piliouras. Gradient Descent Only Converges to Minimizers: Non-Isolated Critical Points and Invariant Regions. In *Innovations in Theoretical Computer Science Conference*, 2017.
- María Pérez-Ortiz, Omar Rivasplata, Emilio Parrado-Hernández, Benjamin Guedj, and John Shawe-Taylor. Progress in Self-Certified Neural Networks. In *NeurIPS 2021 Workshop on Bayesian Deep Learning*, 2021.
- Henning Petzka, Michael Kamp, Linara Adilova, Cristian Sminchisescu, and Mario Boley. Relative Flatness and Generalization. *Conference on Neural Information Processing Systems*, 2021.
- Omar Rivasplata, Ilja Kuzborskij, Csaba Szepesvari, and John Shawe-Taylor. PAC-Bayes Analysis Beyond the Usual Bounds. In *Conference on Neural Information Processing Systems*, 2020.
- André Schlichting. Poincaré and Log-Sobolev Inequalities for Mixtures. *Entropy*, 2019.
- John Shawe-Taylor and Robert C. Williamson. A PAC Analysis of a Bayesian Estimator. In *Conference on Computational Learning Theory*, 1997.
- Jost Springenberg, Alexey Dosovitskiy, Thomas Brox, and Martin Riedmiller. Striving for Simplicity: The All Convolutional Net. In *International Conference on Learning Representations – Workshop Track*, 2015.
- Ilya Tolstikhin and Yevgeny Seldin. PAC-Bayes-Empirical-Bernstein Inequality. In *Conference on Neural Information Processing Systems*, 2013.
- Paul Viallard, Pascal Germain, Amaury Habrard, and Emilie Morvant. A general framework for the practical disintegration of PAC-Bayesian bounds. *Machine Learning*, 2023a.
- Paul Viallard, Maxime Haddouche, Umut Şimşekli, and Benjamin Guedj. Learning via Wasserstein-Based High Probability Generalisation Bounds. In *Conference on Neural Information Processing Systems*, 2023b.
- Paul Viallard, Rémi Emonet, Amaury Habrard, Emilie Morvant, and Valentina Zantedeschi. Leveraging PAC-Bayes Theory and Gibbs Distributions for Generalization Bounds with Complexity Measures. In *International Conference on Artificial Intelligence and Statistics*, 2024a.
- Paul Viallard, Maxime Haddouche, Umut Şimşekli, and Benjamin Guedj. Tighter Generalisation Bounds via Interpolation. *arXiv*, abs/2402.05101, 2024b.
- Cédric Villani. *Optimal transport: old and new*. Springer, 2009.

Kaiyue Wen, Zhiyuan Li, and Tengyu Ma. Sharpness Minimization Algorithms Do Not Only Minimize Sharpness To Achieve Better Generalization. In *Conference on Neural Information Processing Systems*, 2023.

Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-MNIST: a Novel Image Dataset for Benchmarking Machine Learning Algorithms, 2017.

Yun Yue, Jiadi Jiang, Zhiling Ye, Ning Gao, Yongchao Liu, and Ke Zhang. Sharpness-Aware Minimization Revisited: Weighted Sharpness as a Regularization Term. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2023.

Tong Zhang. Information-theoretic upper and lower bounds for statistical estimation. *IEEE Transactions on Information Theory*, 2006.

Appendix A. Supplementary details

A.1. Additional details on Poincaré and Log-Sobolev inequalities

We first provide a proof of Proposition 4.

Proposition 4 *Assume that P is a probability measure on \mathcal{H} such that $dP(h) \propto \exp(-V(h))$ with $V : \mathcal{H} \rightarrow \mathbb{R}$ a smooth function such that $\text{Hess}(V) \succeq \frac{2}{c_{LS}(P)} \text{Id}$.³ Assume that $\ell(h, \mathbf{z}) = \ell_1(h, \mathbf{z}) + \ell_2(h, \mathbf{z})$ with ℓ_1 convex, twice differentiable and ℓ_2 bounded. Then for any $\gamma > 0$, the Gibbs posterior $Q_{\mathcal{S}_m}^\gamma$ is L -Sob(c_{LS}) with constant $c_{LS}(Q_{\mathcal{S}_m}^\gamma) = c_{LS}(P) \exp(4\|\ell_2\|_\infty)$.*

Proof We define P_1 such that $dP_1(h) \propto \exp(-V(h) - \frac{\gamma}{m} \sum_{i=1}^m \ell_1(h, \mathbf{z}_i)) dh$. Then, by the convexity assumption over the loss ℓ_1 , we have $\text{Hess}(V + \frac{\gamma}{m} \sum_{i=1}^m \ell_1(h, \mathbf{z}_i)) \succeq \frac{1}{c_{LS}} \text{Id}$. Then, applying Chafaï (2004, Corollary 2.1), we know that P_1 satisfies a Poincaré inequality with constant $c_{LS}(P)$. Finally, defining P_2 such that $dP_2(h) \propto \exp(-\frac{\gamma}{m} \sum_{i=1}^m \ell_2(h, \mathbf{z}_i)) dh$, thanks to the boundedness of the loss ℓ_2 , we use Guionnet and Zegarliński (2003, Property 4.6), which ensures that $dP_2(h) = Q_{\mathcal{S}_m}^\gamma dP_1(h)$ satisfies a Log-Sobolev inequality with constant $c_{LS}(P) \exp(4\|\ell_2\|_\infty)$. Observing that $P_2 = Q_{\mathcal{S}_m}^\gamma$ completes the proof. ■

For the sake of completeness, we prove Proposition 15, which is Proposition 2.1 of Ledoux (2006), showing that a Log-Sobolev inequality implies a Poincaré inequality.

Proposition 15 (Proposition 2.1 of Ledoux (2006)) *If the distribution Q is L -Sob(c_{LS}), then it is also Poinc(c_P), with $c_P(Q) = \frac{c_{LS}(Q)}{2}$.*

Proof Let $f \in H^1(Q)$ such that $\mathbb{E}_{h \sim Q}[f(h)] = 0$ and $\mathbb{E}_{h \sim Q}[f^2(h)] = 1$. For any $\varepsilon > 0$, we have $1 + \varepsilon f \in H^1(Q)$. We then apply the Log-Sobolev inequality on $1 + \varepsilon f$ to obtain

$$\mathbb{E}_{h \sim Q} [(1 + \varepsilon f(h))^2 (2 \log(1 + \varepsilon f(h)) - \log(1 + \varepsilon^2))] \leq c_{LS}(Q) \varepsilon^2 \mathbb{E}_{h \sim Q} [\|\nabla_h f(h)\|^2].$$

Note that, by a Taylor expansion, we have $\log(1 + \varepsilon f(h)) = \varepsilon f(h) - \frac{(\varepsilon f(h))^2}{2} + o(\varepsilon^2)$ and we have also $\log(1 + \varepsilon^2) = \varepsilon^2 + o(\varepsilon^2)$. Then, plugging this into the previous equation gives

$$\mathbb{E}_{h \sim Q} [2\varepsilon f(h) + 3(\varepsilon f(h))^2 - \varepsilon^2 + o(\varepsilon^2)] \leq c_{LS}(Q) \varepsilon^2 \mathbb{E}_{h \sim Q} [\|\nabla_h f(h)\|^2].$$

We use that $\mathbb{E}_{h \sim Q}[f(h)] = 0$ and we then divide by ε^2 . Taking the limit $\varepsilon \rightarrow 0$ gives:

$$\mathbb{E}_{h \sim Q} [3f(h)^2 - 1] \leq c_{LS}(Q) \mathbb{E}_{h \sim Q} [\|\nabla_h f(h)\|^2].$$

Using that $\mathbb{E}_{h \sim Q}[f^2(h)] = 1$, we obtain

$$1 \leq \frac{c_{LS}(Q)}{2} \mathbb{E}_{h \sim Q} [\|\nabla_h f(h)\|^2].$$

Then, for any $g \in H^1(Q)$ applying this proof on $f : h \mapsto \frac{g(h) - \mathbb{E}_{h \sim Q}[g(h)]}{\sqrt{\text{Var}_{h \sim Q}(g(h))}}$ concludes the proof. ■

3. The notation $A \succeq B$ means that $A - B$ is a semi-definite positive matrix.

A.2. Wasserstein distances

We recall the definitions of the 1-Wasserstein and the 2-Wasserstein distances, which are valid for any predictor space $\mathcal{H} \subseteq \mathbb{R}^d$ equipped with the Euclidean distance.

Definition 16 *The 1-Wasserstein distance between $Q \in \mathcal{M}(\mathcal{H})$ and $P \in \mathcal{M}(\mathcal{H})$ is defined as*

$$W_1(Q, P) = \inf_{\pi \in \Pi(Q, P)} \int_{\mathcal{H}^2} \|x - y\| d\pi(x, y).$$

where $\Pi(Q, P)$ denotes the set of probability measures on \mathcal{H}^2 whose marginals are Q and P . Similarly, the 2-Wasserstein distance between $Q \in \mathcal{M}(\mathcal{H})$ and $P \in \mathcal{M}(\mathcal{H})$ is defined as

$$W_2(Q, P) = \sqrt{\inf_{\pi \in \Pi(Q, P)} \int_{\mathcal{H}^2} \|x - y\|^2 d\pi(x, y)}.$$

A.3. Fundamental difference between time-uniform PAC-Bayes bounds and classical ones

In this work, we establish *time-uniform estimation* PAC-Bayes bounds. Specifically, we focus on the bounds such that there exist $\mathcal{C} \subseteq \mathcal{M}(\mathcal{H})$, a threshold $\varepsilon > 0$, and $\alpha > 0$ such that, with probability at least $1 - \delta$, for all $Q \in \mathcal{C}$ and $m \in \mathbb{N}^*$, we have

$$R_{\mathcal{D}}(Q) \leq \frac{\alpha}{m} [\text{KL}(Q, P) + \log(1/\delta)] + \varepsilon_m.$$

We show that, in a favourable setting, time-uniform estimation PAC-Bayes bounds are sufficient to ensure the almost sure convergence of $(R_{\mathcal{D}}(Q_m))_{m \in \mathbb{N}^*}$ for a posterior sequence $(Q_m)_{m \in \mathbb{N}^*}$, whereas classical PAC-Bayes bounds provide only convergence in probability.

Let us first consider a nonnegative loss ℓ , a prior P , a countable dataset \mathcal{S} , and assume there exists a sequence $(Q_m)_{m \in \mathbb{N}^*}$ which is such that $\text{KL}(Q_m, P) \leq D$ for all $m \in \mathbb{N}^*$ and both $\hat{R}_{\mathcal{S}_m}(Q_m)$ and $\mathbb{E}_{\mathbf{z} \sim \mathcal{D}, h \sim Q_m} [\|\nabla_h \ell(h, \mathbf{z})\|^2]$ go to zero as m goes to infinity almost surely. Then, the classical McAllester's bound states that for any $m \in \mathbb{N}^*$, and $\delta_m \in (0, 1]$, with probability at least $1 - \delta_m$ over \mathcal{S}_m , we have

$$\mathbb{P}_{\mathcal{S}_m} \left(R_{\mathcal{D}}(Q_m) \leq \hat{R}_{\mathcal{S}_m}(Q_m) + \sqrt{\frac{\text{KL}(Q_m, P) + \log(2\sqrt{m}/\delta_m)}{2m}} \right) \geq 1 - \delta_m.$$

Let $\alpha > 0$ and take for all $m \in \mathbb{N}^*$, the confidence parameter $\delta_m = \delta/m$. Since we assume that $\hat{R}_{\mathcal{S}_m}(Q_m) \rightarrow 0$, the square root in McAllester's bound goes to zero. Thus, there exists $m_0 \in \mathbb{N}^*$ such that for all $m \geq m_0$, we have

$$R_{\mathcal{D}}(Q_m) \leq \hat{R}_{\mathcal{S}_m}(Q_m) + \sqrt{\frac{\text{KL}(Q_m, P) + \log(2m\sqrt{m}/\delta)}{2m}} \leq \alpha.$$

Thus, for all $m \geq m_0$, we know that $\mathbb{P}_{\mathcal{S}_m} (R_{\mathcal{D}}(Q_m) < \alpha) \geq 1 - \delta/m$. Taking the limit as m goes to infinity ensures that for any $\alpha > 0$, we have

$$\lim_{m \rightarrow \infty} \mathbb{P}_{\mathcal{S}_m} (R_{\mathcal{D}}(Q_m) < \alpha) = 1.$$

Thus, the classical McAllester’s bound allows the sequence $R_{\mathcal{D}}(Q_m)$ to converge in probability to zero. In contrast, Theorem 6 with $\lambda = 1/C$ is a time-uniform estimation bound with $\alpha = 2C$ and $\varepsilon_m = 2 \left(\hat{R}_{S_m}(Q) + \frac{1}{2C} \mathbb{E}_{\mathbf{z} \sim \mathcal{D}, h \sim Q} [\|\nabla_h \ell(h, \mathbf{z})\|^2] \right)$. Under our assumptions, we know that both $\frac{\alpha}{m} [\text{KL}(Q, P) + \log(1/\delta)]$ and ε_m go to zero as m goes to infinity. Then, taking the limit, we have for all $\delta > 0$

$$\mathbb{P}_{\mathcal{S}} \left(\lim_{m \rightarrow \infty} R_{\mathcal{D}}(Q_m) = 0 \right) \geq 1 - \delta.$$

As $\lim_{m \rightarrow \infty} R_{\mathcal{D}}(Q_m)$ does not depend on δ , we can then make δ go to zero to obtain

$$\mathbb{P}_{\mathcal{S}} \left(\lim_{m \rightarrow \infty} R_{\mathcal{D}}(Q_m) = 0 \right) = 1.$$

Thus, $R_{\mathcal{D}}(Q_m)$ converges almost surely to zero. Such a conclusion cannot be achieved with classical PAC-Bayes bounds; our bound is stronger as it allows an almost sure convergence, which cannot be obtained by simple manipulations of classical bounds.

A.4. A deeper look at the nature of our gradient terms and their originality in PAC-Bayes

Here, we highlight the novelty of incorporating $\nabla_h \ell(h, \mathbf{z})$ in PAC-Bayes bounds and argue that this is not a trivial extension of the PAC-Bayes Bernstein bound. Given a bounded loss $\ell : \mathcal{H} \times \mathcal{Z} \rightarrow [0, 1]$, and starting from the Bernstein bound of Tolstikhin and Seldin (2013), it is possible to derive a bound involving gradient terms with an additional cost when the posterior satisfies a Poincaré inequality. We first restate the PAC-Bayes Bernstein bound of Tolstikhin and Seldin (2013).

Theorem 17 *For any $c_1 > 1$, for any data-free prior $P \in \mathcal{M}(\mathcal{H})$, for any loss function $\ell : \mathcal{H} \times \mathcal{Z} \rightarrow [0, 1]$, and for any $\delta \in (0, 1]$, with probability at least $1 - \delta$ over the sample \mathcal{S} , we have for all distributions $Q \in \mathcal{M}(\mathcal{H})$ that satisfy*

$$\sqrt{\frac{\text{KL}(Q, P) + \ln \frac{\nu_1}{\delta_1}}{(e-2) \mathbb{E}_{h \sim Q} [\text{Var}_{\mathbf{z} \sim \mathcal{D}}(\ell(h, \mathbf{z}))]}} \leq \sqrt{m}, \quad (3)$$

we have

$$R_{\mathcal{D}}(Q) \leq \hat{R}_{S_m}(Q) + (1 + c_1) \sqrt{\frac{(e-2) \mathbb{E}_{h \sim Q} [\text{Var}_{\mathbf{z} \sim \mathcal{D}}(\ell(h, \mathbf{z}))] (\text{KL}(Q, P) + \ln \frac{\nu_1}{\delta})}{m}},$$

where

$$\nu_1 = \left\lceil \frac{1}{\ln c_1} \ln \left(\sqrt{\frac{(e-2)m}{4 \ln(1/\delta)}} \right) \right\rceil + 1.$$

Assuming the technical conditions of Theorem 17, that the distribution \mathcal{D} satisfies a Poincaré inequality, and that the loss $\ell(h, \cdot) \in \text{H}^1(\mathcal{D})$ for all $h \in \mathcal{H}$, we obtain the following corollary.

Corollary 18 *Under the same conditions of Theorem 17, for any distribution \mathcal{D} being Poinc($c_{\mathcal{D}}$), and for any loss $\ell(h, \cdot) \in \text{H}^1(\mathcal{D})$, with probability at least $1 - \delta$ over the sample \mathcal{S} , we have for all distributions $Q \in \mathcal{M}(\mathcal{H})$ satisfying Equation (3)*

$$R_{\mathcal{D}}(Q) \leq \hat{R}_{S_m}(Q) + (1 + c_1) \sqrt{\frac{(e-2) \mathbb{E}_{h \sim Q} \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} [\|\nabla_{\mathbf{z}} \ell(h, \mathbf{z})\|^2] (\text{KL}(Q, P) + \ln \frac{\nu_1}{\delta})}{m}}.$$

Proof From Poincaré inequality (Definition 1), we have

$$\mathbb{E}_{h \sim \mathbb{Q}} \left[\mathbb{V}_{\mathbf{z} \sim \mathcal{D}}(\ell(h, \mathbf{z})) \right] \leq \mathbb{E}_{h \sim \mathbb{Q}} \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} [\|\nabla_{\mathbf{z}} \ell(h, \mathbf{z})\|^2],$$

which is then substituted into the bound of Theorem 17. \blacksquare

By using the additional Poincaré assumption on the data distribution \mathcal{D} , we can obtain a bound with gradient term $\mathbb{E}_{h \sim \mathbb{Q}} \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} [\|\nabla_{\mathbf{z}} \ell(h, \mathbf{z})\|^2]$. However, it remains unclear how $\nabla_{\mathbf{z}} \ell(h, \mathbf{z})$ behaves, as we do not optimise with respect to \mathbf{z} . As a result, these gradients may remain large even after a successful learning phase. In contrast, Theorem 6 involves the term $\mathbb{E}_{\mathbf{z} \sim \mathcal{D}} \mathbb{E}_{h \sim \mathbb{Q}} [\|\nabla_h \ell(h, \mathbf{z})\|^2]$, which is minimised during a successful optimisation process. Moreover, the Poincaré assumption on the data distribution \mathcal{D} is difficult to verify, as we do not know \mathcal{D} . Our Poincaré assumption on the posterior \mathbb{Q} is much easier to verify, as we can choose \mathbb{Q} in practice. A similar discussion then applies for the empirical Bernstein bound (Tolstikhin and Seldin, 2013, Theorem 4) as the empirical variance is with respect to the dataset \mathcal{S} .

A.5. How to make the Wasserstein term go to zero in Theorem 14?

In this section, we prove the following lemma.

Lemma 19 *Assume that the hypothesis set \mathcal{H} is finite and that the loss $\ell : \mathcal{H} \times \mathcal{Z} \rightarrow \mathbb{R}$ satisfies (\star) with constant G . Then, with probability at least $1 - \delta$, the generalisation gap $R_{\mathcal{D}}(h) - \hat{R}_{\mathcal{S}_m}(h)$ satisfies (\star) with the constant $G' = \mathcal{O}\left(G\sqrt{\log(|\mathcal{H}|/\delta)/m}\right)$.*

Proof For the sake of simplicity, let $f : h \mapsto R_{\mathcal{D}}(h) - \hat{R}_{\mathcal{S}_m}(h)$ be the generalisation gap. Then, fix $(h_1, h_2, a, a') \in \mathcal{H}^4$ and our goal is to prove that f satisfies (\star) with another constant G' . First of all, notice that

$$\begin{aligned} & \langle \nabla_h f(h_1) - \nabla_h f(h_2), a - a' \rangle \\ &= \frac{1}{m} \sum_{i=1}^m \langle \nabla_h \ell(h_2, \mathbf{z}_i) - \nabla_h \ell(h_1, \mathbf{z}_i) - (\nabla_h R_{\mathcal{D}}(h_2) - \nabla_h R_{\mathcal{D}}(h_1)), a - a' \rangle. \end{aligned}$$

Moreover, by the condition (\star) for all $\mathbf{z} \in \mathcal{Z}$, we know that

$$|\langle \nabla_h \ell(h_2, \mathbf{z}) - \nabla_h \ell(h_1, \mathbf{z}), a - a' \rangle| \leq G \cdot \|a - a'\| \cdot \|h_1 - h_2\|.$$

Then, by Hoeffding's inequality, applied on the centered random variable $\nabla_h \ell(h_2, \mathbf{z}_i) - \nabla_h \ell(h_1, \mathbf{z}_i) - (\nabla_h R_{\mathcal{D}}(h_2) - \nabla_h R_{\mathcal{D}}(h_1))$ bounded by $G\|a - a'\|\|h_1 - h_2\|$, with probability at least $1 - \delta$, we have

$$|\langle \nabla_h f(h_1) - \nabla_h f(h_2), a - a' \rangle| \leq G \cdot \|a - a'\| \cdot \|h_1 - h_2\| \cdot \sqrt{\frac{2 \log(2/\delta)}{m}}.$$

Taking a union bound on all possible values of $(h_1, h_2, a, a') \in \mathcal{H}^4$ with $\delta' = \delta/|\mathcal{H}|^4$ and a union bound on all tuples yields that, with probability at least $1 - \delta$, for all $(a, a') \in \mathcal{H}^2$,

$$\text{the function } h \mapsto \langle \nabla_h f(h), a - a' \rangle \text{ is } G \sqrt{\frac{2 \log\left(\frac{2|\mathcal{H}|^4}{\delta}\right)}{m}} \|a - a'\| \text{-Lipschitz,}$$

meaning the condition (\star) is verified with constant $G' = G\sqrt{\frac{2\log\left(\frac{2|\mathcal{H}|^4}{\delta}\right)}{m}}$ for the gap. \blacksquare

Appendix B. PAC-Bayes bounds for Lipschitz losses through log-Sobolev inequalities

Extending Catoni's bound to Lipschitz losses. A well-known relaxation of [Catoni \(2007, Theorem 1.2.6\)](#) (see also [Alquier et al., 2016, Theorem 4.1](#)) holding for subgaussian losses has been widely used in practice as a tractable PAC-Bayesian algorithm exhibiting a linear dependency on the KL divergence. Below, we exploit a consequence of the Herbst argument, as stated, for example, in [Ledoux \(2006, Section 2.3\)](#), which asserts that an L -Lipschitz function of a random variable following a distribution \mathcal{D} being L -Sob(c_{LS}) is $L\sqrt{c_{LS}(\mathcal{D})}$ -subgaussian. This leads to the following corollary.

Corollary 20 *For any $\lambda > 0$, for any data-free prior $P \in \mathcal{M}(\mathcal{H})$, for any L -Lipschitz loss $\ell : \mathcal{H} \times \mathcal{Z} \rightarrow \mathbb{R}$ for any $h \in \mathcal{H}$, for any data distribution \mathcal{D} being L -Sob(c_{LS}), with probability at least $1 - \delta$ over \mathcal{S} , for any $Q \in \mathcal{M}(\mathcal{H})$, we have*

$$R_{\mathcal{D}}(Q) \leq \hat{R}_{\mathcal{S}_m}(Q) + \frac{\text{KL}(Q, P) + \log(1/\delta)}{\lambda} + \frac{\lambda^2 L^2 c_{LS}(\mathcal{D})}{2m}.$$

Proof First, we take $f(h, \mathcal{S}_m) := \lambda(R_{\mathcal{D}}(h) - \hat{R}_{\mathcal{S}_m}(h))$ and we use the change of measure inequality ([Csiszár, 1975](#); [Donsker and Varadhan, 1976](#)) to state that, for all $Q \in \mathcal{M}(\mathcal{H})$, we have

$$\mathbb{E}_{h \sim Q} [f(h, \mathcal{S}_m)] \leq \text{KL}(Q, P) + \log \left(\mathbb{E}_{h \sim P} [\exp(f(h, \mathcal{S}_m))] \right).$$

Moreover, Markov's inequality alongside Fubini's theorem gives, with probability at least $1 - \delta$ over the sample \mathcal{S} ,

$$\mathbb{E}_{h \sim Q} [f(h, \mathcal{S}_m)] \leq \text{KL}(Q, P) + \log(1/\delta) + \log \left(\mathbb{E}_{h \sim P} \mathbb{E}_{\mathcal{S}_m} [\exp(f(h, \mathcal{S}_m))] \right).$$

Now, since the loss ℓ is L -Lipschitz on $\mathbf{z} \in \mathcal{Z}$ for all $h \in \mathcal{H}$, we show below that the function f is $\frac{\lambda L}{\sqrt{m}}$ -Lipschitz on the variable \mathcal{S}_m for each $h \in \mathcal{H}$. Indeed, as the loss is L -Lipschitz w.r.t. $\mathbf{z} \in \mathcal{Z}$, for any dataset $\mathcal{S}_m = (\mathbf{z}_1, \dots, \mathbf{z}_m)$, any $\mathcal{S}'_m = (\mathbf{z}'_1, \dots, \mathbf{z}'_m)$, and any $h \in \mathcal{H}$, we have

$$\|f(h, \mathcal{S}_m) - f(h, \mathcal{S}'_m)\| \leq \frac{\lambda L}{m} \sum_{i=1}^m \|\mathbf{z}_i - \mathbf{z}'_i\| = \frac{\lambda L}{m} \sum_{i=1}^m \sqrt{\|\mathbf{z}_i - \mathbf{z}'_i\|^2}.$$

Then, by the concavity of the square root, we have

$$\|f(h, \mathcal{S}_m) - f(h, \mathcal{S}'_m)\| \leq \lambda L \sqrt{\frac{1}{m} \sum_{i=1}^m \|\mathbf{z}_i - \mathbf{z}'_i\|^2} = \lambda L \frac{\|\mathcal{S}_m - \mathcal{S}'_m\|}{\sqrt{m}}.$$

The underlying norm $\|\mathcal{S}'_m\|$ is the one derived from the scalar product $\langle \mathcal{S}_m, \mathcal{S}'_m \rangle = \sum_{i=1}^m \langle \mathbf{z}_i, \mathbf{z}'_i \rangle$. As the distribution \mathcal{D} is L -Sob(c_{LS}), we can deduce that $\mathcal{D}^{\otimes m}$ is also L -Sob(c_{LS}) with the same

constant (Ané et al., 2000, Corollary 3.2.3). Then, using the Herbst argument similarly to Ledoux (2006, Section 2.3), we conclude that for all $h \in \mathcal{H}$, the function $f(h, \cdot)$ is $L\lambda\sqrt{\frac{c_{LS}(\mathcal{D})}{m}}$ -subgaussian. Thus, we have

$$\log \left(\mathbb{E}_{h \sim P} \mathbb{E}_{\mathcal{S}_m} [\exp(f(h, \mathcal{S}_m))] \right) \leq \frac{\lambda^2 L^2 c_{LS}(\mathcal{D})}{2m},$$

which concludes the proof. \blacksquare

Disintegrated PAC-Bayes bounds. Numerical estimation of PAC-Bayes bounds is usually challenging as it often involves Monte-Carlo approximations of the expectation over the posterior Q . A recent line of work (Rivasplata et al., 2020; Haddouche and Guedj, 2022; Viallard et al., 2023a, 2024a) studies *disintegrated PAC-Bayes bounds*, i.e. bounds that hold with high probability on both the dataset \mathcal{S} and a single predictor h drawn from the posterior Q . These bounds may be relevant for practitioners when sampling is easy, as in the case of Gaussian distributions, since they require little computational time. However, a drawback of these bounds is that they do not allow the KL divergence to be used as a complexity measure. Instead, either the disintegrated KL (Rivasplata et al., 2020) or the Rényi divergence (Viallard et al., 2023a) is considered, which can be seen as a relaxation of the KL divergence one. By leveraging the subgaussianity behaviour of Lipschitz losses, it is possible to derive PAC-Bayesian disintegrated bounds, as long as the posterior distribution satisfies a log-Sobolev inequality with a sharp constant (which can be achieved, for instance, for Gaussian distributions with a small operator norm). The new disintegrated bound is introduced in the following lemma.

Lemma 21 *For any L -Lipschitz loss $\ell : \mathcal{H} \times \mathcal{Z} \rightarrow \mathbb{R}$, for any distribution Q being L -Sob(c_{LS}) with $c_{LS}(Q) \leq 1/m$ (and that can depend on the dataset \mathcal{S}_m), with probability $1 - \delta$ over the draw of $h \sim Q$ and \mathcal{S}_m , we have*

$$R_{\mathcal{D}}(h) - \hat{R}_{\mathcal{S}_m}(h) \leq R_{\mathcal{D}}(Q) - \hat{R}_{\mathcal{S}_m}(Q) + \sqrt{\frac{L^2 \log(1/\delta)}{2m}}.$$

Proof We simply remark, by the same argument as in the proof of Corollary 20 that the gap $h \mapsto R_{\mathcal{D}}(h) - R_{\mathcal{S}_m}(h)$ is L -Lipschitz for any \mathcal{S}_m thus the gap is $L\sqrt{c_{LS}(Q)}$ -subgaussian. Then we use

$$R_{\mathcal{D}}(h) - \hat{R}_{\mathcal{S}_m}(h) = \log \left(\exp \left(R_{\mathcal{D}}(h) - \hat{R}_{\mathcal{S}_m}(h) \right) \right).$$

We then apply Markov inequality and exploit the subgaussianness of the gap alongside $c_{LS}(Q) \leq 1/m$ to conclude the proof. \blacksquare

This lemma states that, as long as we assume our loss to be Lipschitz with respect to h , it is possible to easily derive disintegrated PAC-Bayesian bounds. Additionally, Lemma 21 can be easily completed by Corollary 20, which introduces a KL divergence as a complexity term. Note also that, since the loss is Lipschitz, it is also possible to incorporate the 1-Wasserstein distance through the bounds of Haddouche and Guedj (2023b); Viallard et al. (2023b, 2024b). Therefore, having a Log-Sobolev assumption with a sharp constant on the posterior distribution is enough to provide disintegrated PAC-Bayesian bounds involving the KL divergence or the Wasserstein distance, rather than the Rényi divergence or the disintegrated KL divergence.

Appendix C. Proofs

C.1. Proof of Corollary 7

The goal of this section is to prove Corollary 7, which is restated for ease of readability.

Corollary 7 *For any $C > 0$, for any λ such that $\frac{2}{C} > \lambda > 0$, for any data-free prior $P \in \mathcal{M}(\mathcal{H})$, for any loss function $\ell : \mathcal{H} \times \mathcal{Z} \rightarrow \mathbb{R}_+$ such that $\ell(h, \cdot)$ is C^1 almost everywhere on \mathcal{Z} , any $\delta \in (0, 1]$, if the data distribution \mathcal{D} is $\text{Poinc}(c_P)$, with probability at least $1 - \delta$ over the sample \mathcal{S} , for any $m \in \mathbb{N}^*$, any posterior Q being $\text{Poinc}(c_P)$ with $R_{\mathcal{D}}(Q) \leq C$ and such that for any $\mathbf{z} \in \mathcal{Z}$, $\ell(\cdot, \mathbf{z}) \in H^1(Q)$:*

$$R_{\mathcal{D}}(Q) \leq \frac{1}{1 - \frac{\lambda C}{2}} \left(\hat{R}_{\mathcal{S}_m}(Q) + \frac{\text{KL}(Q, P) + \log(1/\delta)}{\lambda m} \right) + \frac{\lambda}{2 - \lambda C} \left(c_P(Q) \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} \left[\mathbb{E}_{h \sim Q} (\|\nabla_h \ell(h, \mathbf{z})\|^2) \right] + c_P(\mathcal{D}) \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} \left(\left\| \mathbb{E}_{h \sim Q} [\nabla_{\mathbf{z}} \ell(h, \mathbf{z})] \right\|^2 \right) \right).$$

To begin this proof, we first state an important intermediate theorem, which holds without any assumptions on the data distribution.

Theorem 22 *For any $C > 0$, for any λ such that $\frac{2}{C} > \lambda > 0$, for any data-free prior $P \in \mathcal{M}(\mathcal{H})$, for any loss function $\ell : \mathcal{H} \times \mathcal{Z} \rightarrow \mathbb{R}_+$, and for any $\delta \in (0, 1]$, with probability at least $1 - \delta$ over the sample \mathcal{S} , for all $m \in \mathbb{N}^*$, for all posterior Q being $\text{Poinc}(c_P)$ with $R_{\mathcal{D}}(Q) \leq C$ and such that for any $\mathbf{z} \in \mathcal{Z}$, $\ell(\cdot, \mathbf{z}) \in H^1(Q)$:*

$$R_{\mathcal{D}}(Q) \leq \frac{1}{1 - \frac{\lambda C}{2}} \left(\hat{R}_{\mathcal{S}_m}(Q) + \frac{\text{KL}(Q, P) + \log(1/\delta)}{\lambda m} \right) + \frac{\lambda}{2 - \lambda C} \left(c_P(Q) \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} \left[\mathbb{E}_{h \sim Q} (\|\nabla_h \ell(h, \mathbf{z})\|^2) \right] + \text{Var}_{\mathbf{z} \sim \mathcal{D}} \left(\mathbb{E}_{h \sim Q} [\ell(h, \mathbf{z})] \right) \right).$$

Theorem 22 highlights the influence of the gradient norm of $\nabla_h \ell(h, \mathbf{z})$ on the generalisation ability: small gradients make the bound vanish. The remaining variance term is not addressed at this stage and can be assumed to be bounded, but we cannot then recover a fast rate.

Proof We start from Chugg et al. (2023, Corollary 17), for any $\lambda > 0$, with probability at least $1 - \delta$, for all $m \in \mathbb{N}^*$, for all posteriors $Q \in \mathcal{H}$, we have

$$R_{\mathcal{D}}(Q) \leq \hat{R}_{\mathcal{S}_m}(Q) + \frac{\text{KL}(Q, P) + \log(1/\delta)}{\lambda m} + \frac{\lambda}{2} \left(\mathbb{E}_{h \sim Q} \left[\mathbb{E}_{\mathbf{z} \sim \mathcal{D}} [\ell(h, \mathbf{z})^2] \right] \right).$$

The last term is then controlled as follows:

$$\mathbb{E}_{h \sim Q} \left[\mathbb{E}_{\mathbf{z} \sim \mathcal{D}} [\ell(h, \mathbf{z})^2] \right] \leq \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} \left[c_P(Q) \mathbb{E}_{h \sim Q} (\|\nabla_h \ell(h, \mathbf{z})\|^2) + \left(\mathbb{E}_{h \sim Q} [\ell(h, \mathbf{z})] \right)^2 \right].$$

We then introduce a supplementary variance term:

$$\begin{aligned}
 &= \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} \left[c_P(\mathbb{Q}) \mathbb{E}_{h \sim \mathbb{Q}} (\|\nabla_h \ell(h, \mathbf{z})\|^2) \right] + \text{Var}_{\mathbf{z} \sim \mathcal{D}} \left(\mathbb{E}_{h \sim \mathbb{Q}} [\ell(h, \mathbf{z})] \right) \\
 &\quad + \left(\mathbb{E}_{\mathbf{z} \sim \mathcal{D}} \mathbb{E}_{h \sim \mathbb{Q}} [\ell(h, \mathbf{z})] \right)^2.
 \end{aligned}$$

Note that by Fubini's theorem, the last term on the right-hand side is exactly $R_{\mathcal{D}}(\mathbb{Q})^2$. Then, using the fact that the averaged true risk is less than C , and reorganising the terms in [Chugg et al. \(2023, Corollary 17\)](#), we obtain, for $\lambda \in (0, \frac{2}{C})$:

$$\begin{aligned}
 R_{\mathcal{D}}(\mathbb{Q}) &\leq \frac{1}{1 - \frac{\lambda C}{2}} \hat{R}_{\mathcal{S}_m}(\mathbb{Q}) + \frac{\text{KL}(\mathbb{Q}, \mathbb{P}) + \log(1/\delta)}{\lambda (1 - \frac{\lambda C}{2}) m} \\
 &\quad + \frac{\lambda}{2 - \lambda C} \left(c_P(\mathbb{Q}) \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} \left[\mathbb{E}_{h \sim \mathbb{Q}} (\|\nabla_h \ell(h, \mathbf{z})\|^2) \right] + \text{Var}_{\mathbf{z} \sim \mathcal{D}} \left(\mathbb{E}_{h \sim \mathbb{Q}} [\ell(h, \mathbf{z})] \right) \right).
 \end{aligned}$$

■

Now that [Theorem 22](#) is proven, we only need to apply the Poincaré assumption on the data distribution to the variance term to derive [Corollary 7](#).

C.2. Proof of [Theorem 8](#)

Theorem 8 *For any $C_1, C_2, c > 0$, for any data-free prior $\mathbb{P} \in \mathcal{M}(\mathcal{H})$, for any loss function $\ell : \mathcal{H} \times \mathcal{Z} \rightarrow \mathbb{R}_+$ being C^2 and for any $\delta \in (0, 1]$, we have, with probability at least $1 - \delta$ over the sample \mathcal{S} , for all $m \in \mathbb{N}^*$, for all \mathbb{Q} being $\text{Poinc}(c_P)=c$, $\text{QSB}(\ell, C_1)$, $\text{QSB}(\|\nabla_h \ell\|^2, C_2)$, and $\ell(\cdot, \mathbf{z}) \in H^1(\mathbb{Q})$, and $\|\nabla_h \ell(\cdot, \mathbf{z})\|^2 \in H^1(\mathbb{Q})$ for all \mathbf{z} ,*

$$\begin{aligned}
 R_{\mathcal{D}}(\mathbb{Q}) &\leq 2\hat{R}_{\mathcal{S}_m}(\mathbb{Q}) + \frac{2c}{C_1} \mathbb{E}_{h \sim \mathbb{Q}} \left[\frac{1}{m} \sum_{i=1}^m \|\nabla_h \ell(h, \mathbf{z}_i)\|^2 \right] \\
 &\quad + 2 \left(C_1 + c \frac{4cG^2 + C_2}{C_1} \right) \frac{\text{KL}(\mathbb{Q}, \mathbb{P}) + \log(2/\delta)}{m}.
 \end{aligned}$$

Proof We start again from [Theorem 6](#), with $\lambda = 1/C_1$, to obtain, with probability at least $1 - \delta/2$:

$$R_{\mathcal{D}}(\mathbb{Q}) \leq 2 \left(\hat{R}_{\mathcal{S}_m}(\mathbb{Q}) + 2C_1 \frac{\text{KL}(\mathbb{Q}, \mathbb{P}) + \log(2/\delta)}{m} \right) + \frac{c_P(\mathbb{Q})}{C_1} \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} \left[\mathbb{E}_{h \sim \mathbb{Q}} (\|\nabla_h \ell(h, \mathbf{z})\|^2) \right]. \quad (4)$$

We now observe that $g : h, \mathbf{z} \mapsto \|\nabla_h \ell(h, \mathbf{z})\|^2$ is nonnegative. Given our assumptions, we apply the proof technique of [Theorem 6](#) on g , *i.e.* we start again from [Corollary 17](#) of ([Chugg et al., 2023](#)), apply Poincaré inequality on \mathbb{Q} and use the QSB assumption on g . We then have, for any $\lambda > 0$, with probability at least $1 - \delta/2$, for all \mathbb{Q} being $\text{Poinc}(c_P)$, $\text{QSB}(g, C_2)$ and $g(\cdot, \mathbf{z}) \in H^1(\mathbb{Q})$ for

all $\mathbf{z} \in \mathcal{Z}$:

$$\begin{aligned} \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} \left[\mathbb{E}_{h \sim \mathcal{Q}} (\|\nabla_h \ell(h, \mathbf{z})\|^2) \right] &\leq \mathbb{E}_{h \sim \mathcal{Q}} \left[\frac{1}{m} \sum_{i=1}^m \|\nabla_h \ell(h, \mathbf{z}_i)\|^2 \right] + \frac{\text{KL}(\mathcal{Q}, \mathcal{P}) + \log(2/\delta)}{\lambda m} \\ &+ \frac{\lambda c_{\mathcal{P}}(\mathcal{Q})}{2} \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} \left[\mathbb{E}_{h \sim \mathcal{Q}} (\|\nabla_h g(h, \mathbf{z})\|^2) \right] + \frac{\lambda C_2}{2} \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} \left[\mathbb{E}_{h \sim \mathcal{Q}} (\|\nabla_h \ell(h, \mathbf{z})\|^2) \right]. \end{aligned} \quad (5)$$

Notice that by definition of $g(\cdot, \mathbf{z}) : \mathbb{R}^d \rightarrow \mathbb{R}$, we have $\nabla_h g(h, \mathbf{z}) = 2\text{Hess}_h(\ell)(h, \mathbf{z})\nabla_h \ell(h, \mathbf{z})$, where $\text{Hess}_h(\ell)$ denotes the Hessian of ℓ . Thus, using the fact that $\ell(\cdot, \mathbf{z})$ is G gradient-Lipschitz for any $\mathbf{z} \in \mathcal{Z}$, we get, for any (h, \mathbf{z}) , that $\|\nabla_h g(h, \mathbf{z})\| \leq 2G\|\nabla_h \ell(h, \mathbf{z})\|$. Substituting this in Equation (5) gives:

$$\begin{aligned} \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} \left[\mathbb{E}_{h \sim \mathcal{Q}} (\|\nabla_h \ell(h, \mathbf{z})\|^2) \right] &\leq \mathbb{E}_{h \sim \mathcal{Q}} \left[\frac{1}{m} \sum_{i=1}^m \|\nabla_h \ell(h, \mathbf{z}_i)\|^2 \right] + \frac{\text{KL}(\mathcal{Q}, \mathcal{P}) + \log(2/\delta)}{\lambda m} \\ &+ \frac{\lambda}{2} (4c_{\mathcal{P}}(\mathcal{Q})G^2 + C_2) \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} \left[\mathbb{E}_{h \sim \mathcal{Q}} (\|\nabla_h \ell(h, \mathbf{z})\|^2) \right]. \end{aligned} \quad (6)$$

Using that $c_{\mathcal{P}}(\mathcal{Q}) = c$, taking $\lambda = \frac{1}{4cG^2 + C_2}$ and reorganising the terms in Equation (6) gives:

$$\begin{aligned} \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} \left[\mathbb{E}_{h \sim \mathcal{Q}} (\|\nabla_h \ell(h, \mathbf{z})\|^2) \right] &\leq 2 \mathbb{E}_{h \sim \mathcal{Q}} \left[\frac{1}{m} \sum_{i=1}^m \|\nabla_h \ell(h, \mathbf{z}_i)\|^2 \right] \\ &+ 2(4cG^2 + C_2) \frac{\text{KL}(\mathcal{Q}, \mathcal{P}) + \log(2/\delta)}{m}. \end{aligned} \quad (7)$$

Finally, taking a union bound and plugging Equation (7) in Equation (4) concludes the proof. \blacksquare

C.3. Proof of Lemma 10

Lemma 10 *For any $\gamma > 0$, for any $m \in \mathbb{N}^*$, for any data-free prior $\mathcal{P} \in \mathcal{M}(\mathcal{H})$ being L -Sob(c_{LS}), for any loss function $\ell : \mathcal{H} \times \mathcal{Z} \rightarrow \mathbb{R}_+$ such that $\ell(\cdot, \mathbf{z}) \in \mathcal{H}^1(\mathcal{P})$ for any \mathbf{z} , we have*

$$\text{KL}(\mathcal{Q}_{\mathcal{S}_m}^\gamma, \mathcal{P}) \leq \frac{\gamma^2 c_{LS}(\mathcal{P})}{4} \mathbb{E}_{h \sim \mathcal{Q}_{\mathcal{S}_m}^\gamma} \left[\|\nabla_h \hat{\mathcal{R}}_{\mathcal{S}_m}(h)\|^2 \right].$$

Proof For conciseness, we rename $\mathcal{Q} := \mathcal{Q}_{\mathcal{S}_m}^\gamma$. We first note that we have

$$\begin{aligned} \text{KL}(\mathcal{Q}_{\mathcal{S}_m}^\gamma, \mathcal{P}) &= \mathbb{E}_{h \sim \mathcal{Q}} \left[\log \left(\frac{d\mathcal{Q}}{d\mathcal{P}}(h) \right) \right] \\ &= \text{Ent}_{\mathcal{P}} \left(\frac{d\mathcal{Q}}{d\mathcal{P}} \right) = \text{Ent}_{\mathcal{P}}[g^2], \end{aligned}$$

where $g = \sqrt{\frac{d\mathcal{Q}}{d\mathcal{P}}}$ and $\frac{d\mathcal{Q}}{d\mathcal{P}}$ is the Radon-Nikodym derivative of \mathcal{Q} with respect to \mathcal{P} . Recall that $\frac{d\mathcal{Q}}{d\mathcal{P}}(h) = \frac{1}{Z} \exp(-\gamma \hat{\mathcal{R}}_{\mathcal{S}_m}(h))$, where $Z = \mathbb{E}_{h \sim \mathcal{P}}[\exp(-\gamma \hat{\mathcal{R}}_{\mathcal{S}_m}(h))]$. Then, the function $g : h \mapsto$

$\frac{1}{\sqrt{Z}} \exp(-\frac{\gamma}{2} \hat{R}_{\mathcal{S}_m}(h))$ belongs to $H^1(\mathbb{P})$ as long as $\ell \in H^1(\mathbb{P})$. Indeed, since \exp is infinitely smooth, $g \in D_1(\mathbb{R}^d)$. Also, as the loss is nonnegative, we have $g \leq \frac{1}{\sqrt{Z}}$, so $g \in L^2(\mathbb{P})$. Finally, we have $\nabla_h g(h) = -\frac{\gamma}{2} g(h) \nabla_h \hat{R}_{\mathcal{S}_m}(h)$. Since $g(h) \leq \frac{1}{\sqrt{Z}}$, we only need to bound $\|\nabla_h \hat{R}_{\mathcal{S}_m}(h)\|^2$ to ensure that $g \in H^1(\mathbb{P})$:

$$\begin{aligned} \|\nabla_h \hat{R}_{\mathcal{S}_m}(h)\|^2 &= \frac{1}{m^2} \sum_{1 \leq i, j \leq m} \langle \nabla_h \ell(h, \mathbf{z}_i), \nabla_h \ell(h, \mathbf{z}_j) \rangle \\ &\leq \frac{1}{2m^2} \sum_{1 \leq i, j \leq m} \|\nabla_h \ell(h, \mathbf{z}_i)\|^2 + \|\nabla_h \ell(h, \mathbf{z}_j)\|^2. \end{aligned}$$

Since we assumed $\|\nabla_h \ell(h, \mathbf{z})\|^2 \in L^2(\mathbb{P})$ for all $\mathbf{z} \in \mathcal{Z}$, we conclude that $g \in H^1(\mathbb{P})$. We then can apply the log-Sobolev inequality to conclude that

$$\begin{aligned} \text{KL}(\mathbb{Q}_{\mathcal{S}_m}^\gamma, \mathbb{P}) &\leq c_{LS}(\mathbb{P}) \mathbb{E}_{h \sim \mathbb{P}} [\|\nabla_h g(h)\|^2] \\ &= \frac{\gamma^2 c_{LS}(\mathbb{P})}{4} \mathbb{E}_{h \sim \mathbb{P}} [\|\nabla_h \hat{R}_{\mathcal{S}_m}(h)\|^2 g^2(h)] \\ &= \frac{\gamma^2 c_{LS}(\mathbb{P})}{4} \mathbb{E}_{h \sim \mathbb{P}} \left[\|\nabla_h \hat{R}_{\mathcal{S}_m}(h)\|^2 \frac{d\mathbb{Q}}{d\mathbb{P}}(h) \right] \\ &= \frac{\gamma^2 c_{LS}(\mathbb{P})}{4} \mathbb{E}_{h \sim \mathbb{Q}_{\mathcal{S}_m}^\gamma} [\|\nabla_h \hat{R}_{\mathcal{S}_m}(h)\|^2]. \end{aligned}$$

■

C.4. Proof of Theorem 11

Theorem 11 We first have $C > 0$, $\gamma > 0$, and a data-free prior $\mathbb{P} \in \mathcal{M}(\mathcal{H})$ being L -Sob(c_{LS}).

(i) For any loss function $\ell : \mathcal{H} \times \mathcal{Z} \rightarrow [0, 1]$ and for any $\delta \in (0, 1]$, with probability at least $1 - \delta$ over the sample \mathcal{S} , for any $m \in \mathbb{N}^*$, we have

$$\mathbb{R}_{\mathcal{D}}(\mathbb{Q}_{\mathcal{S}_m}^\gamma) \leq 2 \left(\hat{R}_{\mathcal{S}_m}(\mathbb{Q}_{\mathcal{S}_m}^\gamma) + \frac{\gamma^2 c_{LS}(\mathbb{P})}{4m} \mathbb{E}_{h \sim \mathbb{Q}_{\mathcal{S}_m}^\gamma} [\|\nabla_h \hat{R}_{\mathcal{S}_m}(h)\|^2] + \frac{\log(1/\delta)}{m} \right).$$

(ii) For any loss function ℓ and prior \mathbb{P} satisfying the conditions of Proposition 4. Then, for any $\frac{2}{C} > \lambda > 0$, with probability at least $1 - \delta$ over the sample \mathcal{S} , for all $m \in \mathbb{N}^*$, assuming $\ell(\cdot, \mathbf{z}) \in H^1(\mathbb{Q}_{\mathcal{S}_m}^\gamma)$ and that $\mathbb{Q}_{\mathcal{S}_m}^\gamma$ is $\mathcal{QSB}(C)$, we have

$$\begin{aligned} \mathbb{R}_{\mathcal{D}}(\mathbb{Q}_{\mathcal{S}_m}^\gamma) &\leq \frac{1}{1 - \frac{\lambda C}{2}} \left(\hat{R}_{\mathcal{S}_m}(\mathbb{Q}_{\mathcal{S}_m}^\gamma) + \frac{\gamma^2 c_{LS}(\mathbb{P})}{4\lambda m} \mathbb{E}_{h \sim \mathbb{Q}_{\mathcal{S}_m}^\gamma} [\|\nabla_h \hat{R}_{\mathcal{S}_m}(h)\|^2] + \frac{\log(1/\delta)}{\lambda m} \right) \\ &\quad + \frac{\lambda e^{4\|\ell_2\|_\infty} c_{LS}(\mathbb{P})}{4 - 2\lambda C} \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} \left[\mathbb{E}_{h \sim \mathbb{Q}_{\mathcal{S}_m}^\gamma} (\|\nabla_h \ell(h, \mathbf{z})\|^2) \right]. \end{aligned}$$

Proof We start again from [Chugg et al. \(2023, Corollary 17\)](#), instantiated with a single λ . Then with probability at least $1 - \delta$, for all posteriors Q and for all $m \in \mathbb{N}^*$, we have

$$R_{\mathcal{D}}(Q) \leq \hat{R}_{\mathcal{S}_m}(Q) + \frac{\text{KL}(Q, P) + \log(1/\delta)}{\lambda m} + \frac{\lambda}{2} \left(\mathbb{E}_{h \sim Q} \left[\mathbb{E}_{\mathbf{z} \sim \mathcal{D}} [\ell(h, \mathbf{z})^2] \right] \right).$$

For the first inequality, we simply take $\lambda = 1$, use the fact that $\ell(h, \mathbf{z})^2 \leq \ell(h, \mathbf{z})$, and reorganise the terms. Finally, we upper-bound the KL term using [Lemma 10](#). For the second inequality, we apply [Proposition 4](#) to use the fact that $Q_{\mathcal{S}_m}^\gamma$ is L -Sob(c_{LS}), and [Proposition 15](#), which ensures that $Q_{\mathcal{S}_m}^\gamma$ is Poinc(c_P), with constant equal to $c_{LS}(Q_{\mathcal{S}_m}^\gamma)/2$. We then follow a proof technique similar to [Theorem 6](#). We have :

$$\mathbb{E}_{h \sim Q_{\mathcal{S}_m}^\gamma} \left[\mathbb{E}_{\mathbf{z} \sim \mathcal{D}} [\ell(h, \mathbf{z})^2] \right] = \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} \left[\text{Var}_{h \sim Q_{\mathcal{S}_m}^\gamma} (\ell(h, \mathbf{z})) + \left(\mathbb{E}_{h \sim Q_{\mathcal{S}_m}^\gamma} [\ell(h, \mathbf{z})] \right)^2 \right].$$

Applying Poincaré inequality then gives:

$$\leq \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} \left[\frac{c_{LS}(P)}{2} e^{4\|\ell_2\|_\infty} \mathbb{E}_{h \sim Q_{\mathcal{S}_m}^\gamma} (\|\nabla_h \ell(h, \mathbf{z})\|^2) + \left(\mathbb{E}_{h \sim Q_{\mathcal{S}_m}^\gamma} [\ell(h, \mathbf{z})] \right)^2 \right].$$

Finally, using the fact that $Q_{\mathcal{S}_m}^\gamma$ is $QSB(\ell, C)$ allow us to reorganise the terms as in [Theorem 6](#). Combining this with [Lemma 10](#) to bound the KL divergence concludes the proof. \blacksquare

C.5. Proof of [Theorem 12](#)

Theorem 12 *For any predictor set \mathcal{H} with finite diameter $D > 0$, and for any function $f : \mathcal{H} \rightarrow \mathbb{R}$ satisfying (\star) , we have for all distributions $P \in \mathcal{M}(\mathcal{H})$ and $Q \in \mathcal{M}(\mathcal{H})$*

$$\mathbb{E}_{h \sim Q} [f(h)] \leq \frac{G}{2} W_2^2(Q, P) + \mathbb{E}_{h \sim P} [f(h)] + D \mathbb{E}_{h \sim Q} [\|\nabla f(h)\|].$$

Proof We first assume that $G = 1$ in the (\star) assumption. We start from the Kantorovich duality formula ([Villani, 2009, Theorem 5.10](#)), instantiated with the cost function $c(x, y) = \|x - y\|^2$. For any Q, P , since W_2 is a distance, we have:

$$W^2(Q, P) = W^2(P, Q) = \sup_{\phi, \psi} \mathbb{E}_{h \sim Q} [\phi(h)] - \mathbb{E}_{h \sim P} [\psi(h)], \quad (8)$$

where the supremum is taken over the functions $\phi, \psi \in L^1(Q) \times L^1(P)$ such that for all $h, h' \in \mathcal{H}^2$, we have $\phi(h) - \psi(h') \leq \|h - h'\|^2$. We claim that if $\phi(h) = f(h) - D\|\nabla f(h)\|$ and $\psi(h') = f(h')$, then the pair (Φ, Ψ) satisfies $\phi(h) - \psi(h') \leq \frac{\|h - h'\|^2}{2}$. Indeed, we have

$$\begin{aligned} \phi(h) - \psi(h') &= f(h) - f(h') - D\|\nabla f(h)\| \\ &= f \circ g(1) - f \circ g(0) - D\|\nabla f(h)\|, \end{aligned}$$

where $g(t) = th + (1 - t)h'$. Then, by the fundamental theorem of calculus, we have

$$\begin{aligned}\phi(h) - \psi(h') &= \int_0^1 (f \circ g)'(t) dt - D \|\nabla f(h)\| \\ &= \int_0^1 \langle \nabla f(th + (1 - t)h'), h - h' \rangle dt - D \|\nabla f(h)\|.\end{aligned}$$

We now control the last term using that $\|h - h'\| \leq D$ and Cauchy-Schwarz inequality:

$$\begin{aligned}\phi(h) - \psi(h') &\leq \int_0^1 \langle \nabla f(th + (1 - t)h'), h - h' \rangle dt - \langle \nabla f(h), h - h' \rangle \\ &= \int_0^1 \langle \nabla f(th + (1 - t)h') - \nabla f(h), h - h' \rangle dt.\end{aligned}$$

Then by the (\star) assumption:

$$\begin{aligned}\phi(h) - \psi(h') &\leq \|h - h'\| \int_0^1 (1 - t) dt \|h - h'\| dt \\ &= \frac{\|h - h'\|^2}{2}.\end{aligned}$$

We then conclude by applying Equation (8) to the pair $(2\phi, 2\psi)$. The general case with $G \neq 1$ follows immediately by considering the pair $(\frac{2}{G}\phi, \frac{2}{G}\psi)$. \blacksquare

C.6. Proof of Corollary 13

Corollary 13 *For any distribution P being L -Sob(c_{LS}) such that $dP(h) \propto \exp(-V(h))dh$, with V being C^2 , for any $R > 0$, for any function f on the centred ball $\mathcal{B}(\mathbf{0}, R)$ of radius R satisfying the (\star) assumption, and for any distribution $Q \in \mathcal{M}(\mathcal{H})$, we have:*

$$\mathbb{E}_{h \sim Q} [f(\mathcal{P}_R(h))] \leq \frac{G c_{LS}(P)}{4} \text{KL}(Q, P) + \mathbb{E}_{h \sim P} [f(\mathcal{P}_R(h))] + 2R \mathbb{E}_{h \sim Q} [\|\nabla_h f(\mathcal{P}_R(h))\|],$$

where \mathcal{P}_R denotes the Euclidean projection onto $\mathcal{B}(\mathbf{0}, R)$.

Proof We fix $R > 0$, and we start from Theorem 12 with predictor space $\mathcal{H}_0 = \mathcal{B}(\mathbf{0}, R)$, where f is gradient-Lipschitz on this ball and the prior and the posterior are respectively $\mathcal{P}_R \# Q$ and $\mathcal{P}_R \# P$. We have

$$\mathbb{E}_{h \sim Q} [f(\mathcal{P}_R(h))] \leq \frac{G}{2} W_2^2(\mathcal{P}_R \# Q, \mathcal{P}_R \# P) + \mathbb{E}_{h \sim P} [f(\mathcal{P}_R(h))] + 2R \mathbb{E}_{h \sim Q} [\|\nabla_h f(\mathcal{P}_R(h))\|].$$

We first prove that $W_2^2(\mathcal{P}_R \# Q, \mathcal{P}_R \# P) \leq W_2^2(Q, P)$. Let $\pi \in \Gamma(Q, P)$ be the optimal transport coupling from P to Q , i.e.

$$W_2^2(Q, P) = \mathbb{E}_{(X, Y) \sim \pi} [\|X - Y\|^2].$$

Then, notice that if we denote by $\pi_1 = (\mathcal{P}_R, \mathcal{P}_R)\#\pi$, then $\pi_1 \in \Gamma(\mathcal{P}_R\#\mathcal{Q}, \mathcal{P}_R\#\mathcal{P})$ and so:

$$\begin{aligned} W_2^2(\mathcal{P}_R\#\mathcal{Q}, \mathcal{P}_R\#\mathcal{P}) &\leq \mathbb{E}_{(X,Y)\sim\pi_1} [\|X - Y\|^2] \\ &= \mathbb{E}_{(X,Y)\sim\pi_1} [\|\mathcal{P}_R(X) - \mathcal{P}_R(Y)\|^2]. \end{aligned}$$

Using the fact that \mathcal{P}_R is 1-Lipschitz gives:

$$\begin{aligned} W_2^2(\mathcal{P}_R\#\mathcal{Q}, \mathcal{P}_R\#\mathcal{P}) &\leq \mathbb{E}_{(X,Y)\sim\pi_1} [\|X - Y\|^2] \\ &= W_2^2(\mathcal{Q}, \mathcal{P}). \end{aligned}$$

Next, we need to control $W_2^2(\mathcal{Q}, \mathcal{P})$. To do so, we use the fact that \mathcal{P} is L -SOB(c_{LS}) to assert, through Otto-Villani's theorem (Otto and Villani, 2000, Theorem 1) that the following holds: $W_2^2(\mathcal{Q}, \mathcal{P}) \leq \frac{c_{LS}(\mathcal{P})}{2} \text{KL}(\mathcal{Q}, \mathcal{P})$. This concludes the proof. \blacksquare

C.7. Proof of Theorem 14

Theorem 14 *Let $\delta \in (0, 1)$ and $P \in \mathcal{M}(\mathcal{H})$ a data-free prior. Assume \mathcal{H} has a finite diameter $D > 0$, for any loss function $\ell : \mathcal{H} \times \mathcal{Z} \rightarrow \mathbb{R}_+$ and any $m \in \mathbb{N}^*$, the generalisation gap $h \mapsto R_{\mathcal{D}}(h) - \hat{R}_{\mathcal{S}_m}(h)$ satisfies (\star) . Assume that $\mathbb{E}_{h\sim P} \mathbb{E}_{\mathbf{z}\sim \mathcal{D}}[\ell(h, \mathbf{z})^2] \leq \sigma^2$, then the following holds with probability at least $1 - \delta$, for any $m > 0$ and any $\mathcal{Q} \in \mathcal{M}(\mathcal{H})$:*

$$R_{\mathcal{D}}(\mathcal{Q}) \leq \hat{R}_{\mathcal{S}_m}(\mathcal{Q}) + \frac{G}{2} W_2^2(\mathcal{Q}, \mathcal{P}) + \sqrt{\frac{2\sigma^2 \log(\frac{1}{\delta})}{m}} + D \mathbb{E}_{h\sim \mathcal{Q}} \left(\left\| \nabla_h R_{\mathcal{D}}(h) - \nabla_h \hat{R}_{\mathcal{S}_m}(h) \right\| \right).$$

Proof We start from Theorem 12, using the fact that $R_{\mathcal{D}}(h) - \hat{R}_{\mathcal{S}_m}(h)$ is G -gradient-Lipschitz for any $m \in \mathbb{N}^*$ to obtain:

$$\begin{aligned} \mathbb{E}_{h\sim \mathcal{Q}} [R_{\mathcal{D}}(h) - \hat{R}_{\mathcal{S}_m}(h)] &\leq \frac{G}{2} W_2^2(\mathcal{Q}, \mathcal{P}) + \mathbb{E}_{h\sim \mathcal{P}} [R_{\mathcal{D}}(h) - \hat{R}_{\mathcal{S}_m}(h)] \\ &\quad + D \mathbb{E}_{h\sim \mathcal{Q}} \left(\left\| \nabla_h R_{\mathcal{D}}(h) - \nabla_h \hat{R}_{\mathcal{S}_m}(h) \right\| \right) \end{aligned}$$

The only remaining term to control is $\mathbb{E}_{h\sim \mathcal{P}} [R_{\mathcal{D}}(h) - \hat{R}_{\mathcal{S}_m}(h)]$. For this, we use the supermartingale concentration inequality of Chugg et al. (2023, Corollary 17) instantiated with the prior equals to the posterior, which shows that, for any $\lambda > 0$, with probability at least $1 - \delta$, we have

$$\mathbb{E}_{h\sim \mathcal{P}} [R_{\mathcal{D}}(h) - \hat{R}_{\mathcal{S}_m}(h)] \leq \frac{\log(1/\delta)}{\lambda} + \frac{\lambda}{2} \mathbb{E}_{h\sim \mathcal{P}} \mathbb{E}_{\mathbf{z}\sim \mathcal{D}} [\ell(h, \mathbf{z})^2].$$

The last term on the right-hand side is bounded by σ^2 by assumption. Taking $\lambda = \sqrt{\frac{2\log(1/\delta)}{\sigma^2}}$, we finally get $\mathbb{E}_{h\sim \mathcal{P}} \mathbb{E}_{\mathbf{z}\sim \mathcal{D}} [\ell(h, \mathbf{z})^2] \leq \sqrt{2\log(1/\delta)/m}$, which concludes the proof. \blacksquare